From self-enforcing agreements to self-sustaining norms

Mikhail Safronov
University of Cambridge

Bruno Strulovici*
Northwestern University

December 23, 2017

Abstract

Self-enforcing agreements, such as relational contracts and international agreements, prescribe actions and rules enforcing these actions, but typically ignore agents' ability to change these rules. This paper studies self-sustaining norms, which prescribe how individuals react to one another's actions but also to proposals to change the rules. We characterize the set of self-sustaining norms when agents interact frequently, which has a remarkably simple structure. Inefficient norms may arise even when all actions are public and frequent and agents can credibly commit to rules.

1 Introduction

Economists often distinguish agreements that rely on commitment from those that are self-enforcing. Relational contracts (MacLeod and Malcomson (1989), Levin (2003)), sovereign debt contracts (Bulow and Rogoff (1989, 1991), Atkeson (1991)), dynamic provision of public goods (Levhari and Mirman (1980), Fershtman and Nitzan (1991)) and resolutions of the tragedy of the commons, and cooperative equilibria of repeated games are well-known instances of self-enforcing agreements, each of which represents an important paradigm in its respective field of study. But are these agreements really immune from any commitment assumption?

In general, an agreement is viewed as self-enforcing if violations can be deterred by punishments which are continuation equilibria of the game underlying the agreement. As seminal studies of renegotiation-proof equilibria have noted, this concept ignores the possibility, for the parties involved, of moving away from the punishment continuation to another continuation that all par-

*We are grateful for comments from David Miller, Larry Samuelson, Ilya Segal, Tadashi Sekiguchi, Takuo Sugaya, Joel Watson, and seminar participants at the 2015 SAET meeting, the 2015 World Congress of the Econometric Society, and the 2016 Asian Meeting of the Econometric Society. Strulovici gratefully acknowledges financial support from an NSF CAREER Award (Grant No. 1151410) and a fellowship from the Alfred P. Sloan Foundation

ties prefer to the punishment.[1] In this scenario, what prevents the parties from moving to a Pareto-superior continuation? If such a move is feasible, doesn't this compromise the relevance of "self-enforcing" equilibria? And if it isn't, doesn't this, in fact, constitute a form of commitment?

These questions suggest a stronger notion of self-enforcement, according to which an agreement must deter not only deviations in actions, but also proposals which undermine the punishments necessary to enforce the initial agreement. This stronger notion raises several conceptual challenges, which this paper aims to identify and resolve.

The first challenge is a modeling question: to consider Pareto-improving proposals, one must envisage a larger game in which such proposals are entertained. Should this larger game be implicit, as in the large literature on renegotiation-proof contracts?[2] Or should it be formalized explicitly by a model in which continuation equilibria can be negotiated? Put differently, should renegotiation be considered as a possibly on-path phenomenon of a larger game? Or can it be conceptualized without loss of generality as a constraint imposed on equilibria of the underlying game?

The cooperative approach to renegotiation–by far the most common in repeated games[3]—has focused on defining concepts of renegotiation-proof equilibrium. In some settings, however, on-path

[1]This point is made explicitly by Farrell (1983) and Abreu and Pearce (1991) and has also been noted in the more applied literature on self-enforcing agreements. For example, Levin (2003) shows that the efficient contract that he characterizes is "strongly" optimal, i.e., efficient after all histories. In the context of sovereign-debt contracts, Kletzer and Wright (2000) identify a self-enforcing agreement that is renegotiation-proof according to various definitions of the concept. In these works, the issue of renegotiation-proofness essentially collapses due to the availability of large transfers and quasi-linear utility. By contrast, in a version of Levhari and Mirman's "fish war" model of dynamic provision of public goods, Cave (1987) identifies several non trivial issues with renegotiation—such as the nonexistence of strong perfect equilibria.

[2]Early work following this approach include Farrell (1983), Pearce (1987), DeMarzo (1988), Bernheim and Ray (1989), Farrell and Maskin (1989), Abreu and Pearce (1991), Asheim (1991), Bergin and MacLeod (1991), Abreu, Pearce, and Stacchetti (1993), Benoît and Krishna (1993) and Ray (1994). While some of this literature is motivated by a critique of self-enforcing agreements, viewed as equilibria of a dynamic game which may not survive renegotiation, the remedies proposed by this literature all take the form of restrictions—usually, based on efficiency considerations—on the set of equilibria and use a cooperative or axiomatic approach. By contrast, in the present paper, a self-sustaining norm may be viewed as a self-enforcing agreement in a larger game that includes renegotiation. Blume (1994) stands out in this literature as it alludes to a larger game to justify efficient (or "renegotiation-perfect") bargaining outcomes. The actual process of renegotiation is unmodeled and captured instead by an efficient bargaining solution concept.

[3]An important exception is Miller and Watson (2013), who study bargaining in repeated games with transfers. A key axiom in their theory is the concept of "no-fault disagreement" (NFD), which stipulates that the continuation of the game when bargaining breaks down in a given period or when a player fails to make a prescribed transfer in that period be independent of what triggered this event. The NFD axiom together with the availability of large transfers as well as internal and external efficiency axioms weaker than those typically invoked under the cooperative approach allows them to provide a sharp characterization of bargaining outcomes in their environment. In Appendix L, we explore the consequences of their NFD axiom in our theory while maintaining the absence of transfers.

renegotiation allows outcomes that cannot be replicated by renegotiation-proof equilibria. This is for example the case when renegotiation entails some frictions (Brennan and Watson (2013)). To determine whether the focus on renegotiation-proof equilibria is without loss, and which notion of renegotiation-proofness to use, if any, this paper incorporates renegotiation stages into an underlying model of dynamic interactions.[4]

We consider a simple environment, in which each period of a repeated game is enlarged to include a renegotiation stage: a player is chosen at random to propose a new continuation of the enlarged game, which may be accepted or rejected by other players. To achieve maximal clarity and focus on the core conceptual issues, our benchmark setting concerns two players and a renegotiation stage that involves only one proposal by one of the players, followed by an acceptance decision by the other player. Equilibria in this enlarged game are called *norms* to emphasize that, unlike equilibria of the underlying game, they prescribe not only how agents should react to each other's actions, but also which proposals they should make and how they should react to any proposal.[5] The norms that we consider are thus similar to the *cultural beliefs* studied by Greif (1994) in that they include both the rules governing actions but also a deeper characteristic concerning how individuals perceive proposals to change the rules and react to them.

To give a substantive meaning to proposals, we say that a norm is *self-sustaining* if it satisfies a simple equilibrium refinement: any proposal that is accepted is played.[6] Self-sustainability is thus a property that a norm may have, a rule dictating how to handle incoming proposals. Of course, other rules exist—such as rules stipulating that all proposals be ignored—and players are allowed to propose, and possibly implement, any change to these rules. We characterize the set of self-sustaining norms as the discount factor goes to 1, and then show that i) more complex renegotiation protocols do not affect our results, ii) our refinement has an equivalent set-theoretic formulation which differs from existing ones, and iii) the analysis and results extend to more players, but raise new conceptual issues.

[4]Some papers have considered a "renegotiation-proofness" principle in specific environments (e.g., Hart and Tirole (1988)). It should be clear that this approach is limited in dynamic games with an infinite horizon, most obviously because there is no agreed upon concept of what "renegotiation-proof" means.

[5]This terminology should not be confused with the "social norms" considered by DeMarzo (1990) or Asheim (1991), which are subsets of equilibria. By modeling norms as equilibria of a dynamic game, we address Ray's (1994) insight that renegotiation-proof sets should be recursively consistent, but instead of applying it to sets of equilibria in the underlying game, we apply it to the enlarged game with renegotiation, by requiring that our norms be perfect equilibria of the enlarged game.

[6]More precisely, the refinement is that any accepted proposal including, possibly, the first off-path proposal, be played. The terminology refers to the fact that the norm must survive challenges from proposals to change it, while treating all proposals seriously.

When renegotiation is explicitly modeled as part of a larger dynamic game, new possibilities emerge which were occulted by the cooperative approach to renegotiation. Consider, in a two-player repeated game, a continuation equilibrium that is Pareto dominated by another continuation. If someone proposes to switch to the Pareto-improving continuation, the proposal's recipient should consider the consequences of rejecting the proposal, which need not coincide with the continuation equilibrium in the absence of any proposal: just as in any dynamic game, continuations can a priori depend on all histories.

To be concrete, consider a repeated game and history with three possible continuation equilibria, s_1, s_2, and s_A, such that s_1 is Pareto dominated by s_2, and s_A is most preferred by Ann but least preferred by Bob. If s_1 was to be played but Bob could, in larger game with negotiation, suggest s_2 instead, there is no reason why Ann should accept Bob's proposal if rejecting it leads to s_A. Similarly, Bob could be rewarded for rejecting Ann's proposal to move to s_2 by a fourth equilibrium s_B that is most preferred by Bob and least preferred by Ann. In this situation, s_1 is stable despite being Pareto dominated, because all proposals are deterred by punishing the proposer and rewarding the rejector of the proposal.

This observation suggests a novel mechanism for the persistence of inefficient outcomes in environments with dynamic interactions: even when agents i) interact arbitrarily frequently; ii) can perfectly observe one anothers' actions, iii) communicate freely; and iv) can credibly agree to switch to more efficient equilibria, inefficient outcomes may still be sustained by a norm which deters proposals by rewarding rejectors, as in the example above. This mechanism, which has applications to political economy, regulatory policy, and other fields, may be viewed as a social norm which prescribes players' reaction to deviating actions but also to deviating proposals.

This mechanism is consistent with the impact on social interactions and institutions of Greif's (1994) cultural beliefs, which our self-sustaining norms seem to capture: Greif notes in particular (p. 925) that "even if each member of the society recognizes the inefficiency caused by individualist cultural beliefs, a unilateral move by an individual or a (relatively) small group would not induce a change. Expectations about expectations are difficult to alter, and thus cultural beliefs can make Pareto-inferior institutions and outcomes self-enforcing."

One of the hardest conceptual questions here is to understand which proposals can be deterred by a norm. Suppose that Ann and Bob are stuck at s_1, as previously described. Could they agree to tear down the rule which stipulates that rejecting s_2 leads to s_A or s_B? But proposing such a rule-tearing agreement amounts to making a proposal, not unlike the proposal to replace s_1 by s_2. In the enlarged game, players anticipate all possible proposals, and all proposals are subject to the argument presented earlier.[7]

These observations point to a novel formal definition of social norms that is dynamic and encompasses the consideration of other norms. Some authors (many considering cooperative or set-theoretic concepts of renegotiation) have described a norm as a set of equilibria which players view as "acceptable" or "credible." Thus, if a continuation equilibrium is Pareto dominated by a second one, this domination is not a concern if the second equilibrium is not in the norm, i.e., not credible. However, one may alternatively view a norm as the equilibrium that is actually played on path. For example, driving on the right side of the road is the "norm" in the United States, but driving on the left side of the road is the norm in Japan. Should a norm be conceptualized as a set of acceptable equilibria, as with the cooperative approach, or as the actual equilibrium that players are engaged in (which may be a layman's view of the concept)? A related question is whether a norm should be time-independent, as in most papers following the cooperative approach, or allowed to vary over time? Some authors, notably Abreu and Pearce (1991) and Asheim (1991) have argued that norms may vary over time, and noted that even presumably weak concepts of renegotiation-proofness fail to allow for this.[8]

Our framework answers both questions simultaneously, by defining a norm as an equilibrium of the dynamic game enlarged to include proposals: it specifies not only what is played on path (e.g., driving on the right side of the road), but also how players should consider proposals to change the norm, i.e., which proposals are considered acceptable, and when. In particular, a proposal which may have been acceptable at earlier stages of the game may become unacceptable after some histories, such as those that include a recent deviation.

While some credible Pareto-improving proposals can be deterred, as explained above, this does not mean that renegotiation has no bite. Consider, for example, a pure coordination game in which the

[7] A related question concerns how a proposal should be interpreted. In this paper, we define proposals as messages, as in Miller and Watson (2013), each of which is associated with an equilibrium of the enlarged game. One could alternatively imagine a hierarchy of proposals, in which each proposal of level $k+1$ dictates the rule concerning proposals of level k. Unless an artificial cap is imposed on this hierarchy (which would amount to a form of commitment), this suggests the consideration of a universal proposal space in which each proposal includes a prescription on how to react to any proposal from this space. Interestingly, however, this construction is impossible: no set of messages can be rich enough to include as distinct proposals all equilibria that it generates. The set of equilibria depends on the set of messages, since each equilibrium prescribes a reaction to each message. This leads to recursive equation which puts the set of equilibria in bijection with its power set. This equation has no solution, by Cantor's impossibility theorem (see, e.g., Mendelson (1997)). The argument is explained further in Appendix D.

[8] While Abreu and Pearce (1991) allow the set of credible deviations to differ from the set of on-path equilibria, they impose that the set of credible deviations always be the same. We do not impose this restriction here, although a payoff-equivalent formulation of self-sustaining norms, described in Section 5, has a similar structure.

Pareto frontier consists of a single payoff vector. Then, any proposal to move to this point would be accepted, since rejecting it could only lead to a worse outcome for all players.

We characterize the set of all self-sustaining-norm payoffs as the discount factor goes to 1, and show that this set largely depends on the nature of the underlying game. In games with high conflict, such as the prisoners' dilemma, renegotiation has no impact on the Folk Theorem: every individually rational payoff can be implemented by a self-sustaining norm. In more cooperative games, renegotiation eliminates many inefficient outcomes. In general, our results provide a way of assessing the amount of conflict or misalignment in a game, providing a continuous spectrum for the impact of renegotiation. It turns out that the payoffs implementable by self-sustaining norms have a simple characterization: they coincide with the set of all individually rational payoff vectors that lie above a particular point whose location depends on the shape of the Pareto frontier. Except in pure coordination games, there exist inefficient self-sustaining norms.

In a political or regulatory context, these inefficient norms may be interpreted as rules set by a third party who benefits from the inefficiency between the players. For example, in a dictatorship, rebellions may be deterred by rules stipulating that anyone proposing a rebellion sees his proposals rejected and is punished by other players who are rewarded for punishing him. While all citizens prefer the rebellion over the status quo, the rules guarantee that it is in no one's interest to ever proposal a rebellion, even when all actions are public and an agreed proposal would be binding.

With an arbitrary number of agents, several new issues emerge. First, could a subset of agents make a partial agreement, possibly at the expense of other agents? What would this imply for other agents' strategies? Could such agreements be private and gradually discovered by the remaining agents? While these issues raise important and interesting challenges,[9] we focus here on global agreements: a proposal is submitted publicly to the entire set of players, and must be approved by a supermajority of the players. As in any model of supermajority vote, the outcome is binary: either the proposal passes, or it loses and a different continuation equilibrium ensues. As in the two-player case, this continuation can, however, depend on the identity of the proposer.[10] A norm

[9]When agreements are public, these questions relate to the study of coalition formation and coalition-proof equilibria (Bernheim, Peleg, and Whinston (1987)). And even if one can give a meaningful private agreements, such agreements may generate asymmetric information among players and create well-known challenges to analyze the formation and behavior of subsequent coalitions, as pointed out by Crémer (1996) and in the context of auctions, by McAfee and McMillan (1992), and Caillaud and Jehiel (1994). See also Che and Kim (2006) for a more recent treatment of this issue.

[10]Our approach is consistent with a large literature in political economy and social choice, according to which the implications of a vote depend only on whether a winning coalition was achieved, not on the composition of the winning coalition. Allowing the continuation equilibrium to depend finely on the identity of individuals who voted in favor of the proposal would de factor split any proposal into a complex proposal *schedule* describing all possible continuations,

is self-sustaining if it satisfies the refinement that any proposal accepted by at least L agents, where L is the supermajority threshold, is played. Our main result extends to any number of players: for any supermajority rule, the set of payoffs implemented by self-sustaining norms consists of all payoffs that lie above a particular payoff vector. Moreover, this set becomes larger (i.e., more permissive) as the supermajority rule becomes more stringent (i.e., L becomes larger).

2 Setting

We consider a repeated game with renegotiation, in which each period features an action stage followed by a renegotiation stage. For expositional simplicity, the benchmark setting concerns two players and focuses on the simplest renegotiation protocol: in each period, one player (at most) gets to make a proposal, which may be accepted or rejected before moving to the next period. Section 6 extends the analysis to more agents and Appendix G considers more general protocols of proposals, showing that the main results are unaffected by this generalization.

For $i \in \{1, 2\}$, player i's stage-game action, a_i, lies in a finite set denoted $\mathcal{A}_i$. The vector $\mathbf{a} = (a_1, a_2)$ determines current-period payoffs $\mathbf{u}(\mathbf{a}) = (u_1(\mathbf{a}), u_2(\mathbf{a}))$. A distribution α_i over $\mathcal{A}_i$ is a *mixed action* for i, and $\alpha = (\alpha_1, \alpha_2)$ denotes the vector of mixed actions for both players. Players put a weight $\varepsilon \in (0, 1)$ on the current period, which corresponds to a common discount factor $\delta = 1 - \varepsilon$.

Each period consists of the following stages:

1) Players observe the realization z of a public randomization device taking values in $[0, 1]$;

2) They simultaneously and privately choose a mixed strategy $\alpha_i \in \Delta(\mathcal{A}_i)$, $i \in \{1, 2\}$.[11] Condi-

and raise a number of strategic considerations (such as the potential benefit of voting against a proposal which one knows will pass anyway), which would require a separate analysis. We do show that if the continuation equilibrium can depend arbitrarily on the identity of the rejectors, then the Folk Theorem can be restored, see Appendix J.

[11] In accordance with current practice, we allow players to use privately mixed strategies. This feature distinguishes our analysis from some of the earlier work on renegotiation. For example, Farrell and Maskin (1989) assume that players can observe each other's mixing strategies, rather than just the realized actions. This distinction can severely affect the set of weakly renegotiation-proof (WRP) equilibria, the concept introduced by Farrell and Maskin. Appendix K provides an example in which all Pareto efficient WRP (which are known to exist) are destroyed and the repetition of an inefficient stage-game Nash equilibrium is the only WRP when mixing is private. Intuitively, when players observe each other's mixed strategy, there is without loss a single continuation payoff vector, conditional on players' mixed strategies. When mixtures are unobservable, however, there must be a continuation vector for each possible outcome of the mixture, chosen so as to make each player indifferent across all actions in the support of his mixed strategy. Moreover, all of these vectors must belong to the renegotiation-proof set. This is problematic because some of these continuations may have Pareto-ranked payoffs, violating weak renegotiation-proofness. Bernheim and

tional on the realization z of the public randomization device, players choose their mixed actions independently from each other;

3) The vector $\mathbf{a}$ of actions is observed and the period's payoffs are realized;

4) With probability $p < 1$, one of the players is given an opportunity to send a *message* from a set $\mathcal{M}$, whose cardinality weakly exceeds the cardinality of the continuum. Each player has the same probability of $\frac{p}{2}$ being chosen.[12] The chosen player may conceal his opportunity to send a message by remaining silent, or mix between sending a message or staying silent;

5) If i sent a message, player $-i$ decides whether to accept it, possibly mixing between acceptance and rejection. The resulting decision D_{-i} equals 1 if $-i$ accepts the message and 0 if he rejects it;[13]

The public history of a period consists of a realisation z of the randomization device; an action vector $\mathbf{a}$; a (possibly empty) message m_i by one of the players, denoted i, and if i sent a nonempty message m_i, an acceptance decision D_{-i}. In addition, each player privately observes the mixing probability used for each of his decisions.

We focus on public equilibria of this game, which we will sometimes call ***norms*** to indicate that they concern the enlarged game rather, and let $\mathcal{S}$ denote the set of all norms. Since messages can always be ignored regardless of whether they were accepted, the set of players' expected payoffs across all norms contains the set $V(\delta)$ of continuation payoffs of the underlying repeated game without renegotiation. Moreover, the reverse inclusion still holds, as any norm in $\mathcal{S}$ can be replicated without renegotiation: During the renegotiation stage players' behavior may affect their continuation payoffs in the next period. Before the renegotiation stage, players consider these continuation payoffs as random variables, which depend on which player gets a chance to propose, and (mixed) equilibrium strategies. Since the cardinality of players' continuation payoffs coincides with the cardinality of outcomes of the public randomization device, the randomization device can simulate the outcome of renegotiation, yielding the same distribution over continuation payoffs in the next period. This proves the following lemma.

Lemma 1 *The set of continuation payoffs implemented by all the norms $\mathcal{S}$ coincides with $V(\delta)$.*

Our main concept is an equilibrium refinement applied to norms.

Ray (1989) rule out mixing altogether, focusing their analysis on pure-strategy equilibria.

[12]When players have different probabilities of making a proposal, the sufficient conditions are unchanged and the necessary conditions entail a payoff lower bound on each player that increases with this player's proposal probability, consistent with the intuition that a higher proposal probability means an increased bargaining power. See Appendix E.

[13]The paper's main results hold as stated when the renegotiation stage includes multi rounds of negotiation. See Appendix G.

Let $\mathcal{H}^+$ denote the set of all finite public histories ending after an action stage such that no off-path proposal has been accepted.[14]

Definition 1 *A norm is* self-sustaining *if the following holds for any history $h \in \mathcal{H}^+$:*

1. **Message Richness** *Each message $m \in \mathcal{M}$ is assigned a norm of $\mathcal{S}$, and the payoff vectors spanned by these norms cover all of $V(\delta)$;*[15]
2. **Binding Acceptance** *If a message is proposed and accepted immediately following h, the norm assigned to it is played from the next period onward.*

Definition 1 thus requires that all accepted on-path proposals as well as the first accepted off-path proposal be implemented. Put differently, Definition 1 means that i) as long as players have not rejected it, they obey the rule stipulating that any accepted proposal is played, ii) the only way of departing from rule is for them to make and accept an off-path proposal that violates the rule.[16]

Under a self-sustaining norm, players can credibly propose, accept, and thus implement, any alternative norm. However, a player on the receiving end of any Pareto-improving off-path proposal finds it optimal to reject it. We emphasize that, under a self-sustaining norm, any accepted proposal is binding even when the proposed itself is not self-sustaining. Indeed, an accepted proposal governs not only how players respond to each others' actions, but also how they respond to each others' proposals. It is perfectly admissible, for instance, for players to agree at some point to ignore all subsequent proposals—regardless of the response to these proposals—just as they may agree to ignore specific deviations in the underlying game.

Self-sustaining norms have a set-theoretic formulation provided in Section 5: we define a *convention* as a set of norms, and introduce a notion of stability for conventions. A norm is self-sustaining if and only if the convention consisting of all continuations of the norm—following histories at which no off-path proposal was accepted—is stable, with a reciprocal statement provided in Section 5. We also introduce a notion of "credible" proposals[17] and show that restricting the self-sustainability

[14] Formally, let $\mathcal{H}$ denote the set of all finite public histories ending after an action stage: elements of $\mathcal{H}$ take the form $h = (z_0, \mathbf{a_0}, \mathbf{m_{0,i_0}}, \mathbf{D}_{\mathbf{0},-\mathbf{i_0}}, \ldots, \mathbf{z_t}, \mathbf{a_t}, \mathbf{m_{t,i_t}}, \mathbf{D}_{\mathbf{t},-\mathbf{i_t}}, \mathbf{z_{t+1}}, \mathbf{a_{t+1}})$ for some $t+1 \in \mathbb{N}$, where $m_{t,i_t} = D_{t,-i_t} = \emptyset$ when no one proposed in period t. Then $\mathcal{H}^+$ consists of all histories $h \in H$ with the following property: for any period t' covered by h for which that $m = m_{t',i_{t'}}$ is nonempty, either m is in the support of $i_{t'}$'s on-path proposals, given the history h truncated after the action stage of period t', or $D_{t',-i_{t'}} = 0$, i.e., $-i_{t'}$ rejects m.

[15] This condition is always achievable: By assumption, $\mathcal{M}$ has the cardinality of the continuum. Therefore, each payoff vector v of $V(\delta)$ can be mapped to some message m, by assigning to m a norm that implements v.

[16] Of course, off-path proposals do not necessarily imply a rejection of this rule.

[17] A proposal is *credible* relative to a convention of norms if any ulterior deviation from the proposal, whether at the action or the proposal stage, triggers a continuation norm in the convention.

refinement to "credible" proposals yields the same necessary and sufficient conditions as those obtained when all proposals are included in the refinement.

3 Main Result

Our main objective is to study the set of payoffs achieved by self-sustaining norms as players' discount factor δ converges to 1, for each level of renegotiation frictions.

To achieve this objective, we adjust the probability p of a renegotiation opportunity in each period proportionally to the period's weight $\varepsilon = 1-\delta$. This normalization has the following interpretation: if ε represents the duration of each period and $p = q\varepsilon$ for some parameter $q > 0$, the probability that a proposal opportunity arises within one unit of time, i.e., within $K = 1/\varepsilon$ periods, is $1-(1-p)^K = 1-(1-q\varepsilon)^{1/\varepsilon} \sim q$. The parameter q thus stands for the *frequency* of proposal opportunities per unit of time. This normalization is useful to capture a non-degenerate effect of renegotiation frictions.[18]

Given $q \geq 0$, a payoff vector X is *sustainable at frequency* q if there is a threshold $\varepsilon(X) > 0$ such that for all $\varepsilon \in (0, \varepsilon(X))$, X is implemented by some self-sustaining norm associated with discount factor $\delta = 1-\varepsilon$ and renegotiation frequency q. We let $V(q)$ denote the set of all payoff vectors which are sustainable at frequency q and $\bar{V} = \cap_{q\geq 0} V(q)$. A payoff vector X belonging to $\bar{V}$ is implementable regardless of the renegotiation frequency and is said to be *sustainable*.

Let $\underline{v}_i$ denote i's minmax payoff in the stage game[19] and P_i denote the feasible payoff vector that gives i his maximal payoff among all payoff vectors above the minmax. If several such vectors exist, the vector whose payoff for $-i$ is the lowest is chosen. The weak individually-rational Pareto

[18] Theorem 1 implies that if the probability p of per-period renegotiation opportunity is fixed independently of ε, the necessary condition for self-sustainability reduces to the sufficient condition described by the theorem, which is independent of p. With p independent of $\varepsilon = 1-\delta$, taking δ to 1 essentially eliminates any renegotiation friction, regardless of p, as any failure to renegotiate an inefficient norm in the current period will be followed by an arbitrarily close opportunity to renegotiate. Moreover, the exploding renegotiation frequency which results from this assumption also creates an instability as δ goes to 1: For each $p > 0$, there are examples for which self-sustaining norms fail to exist as δ goes to 1, described in Appendix I. Our example features an inefficient stage-game Nash equilibrium as well as Pareto-efficient equilibria which must be sustained by the threat of large punishments. When renegotiation opportunities arrive at fixed rate $p > 0$ per period and δ goes to 1, the expected length of punishments becomes too short for the Pareto efficient vectors above the stage-game Nash equilibrium to be implementable—the gain from a deviation in action is proportional to $\varepsilon = 1-\delta$ (the weight put on the current period), while the duration of any punishment is of order $p\varepsilon$ (i.e., until it is renegotiated) and thus smaller than the deviation gain, for judiciously chosen payoffs. The inefficient stage-game Nash equilibrium is not self-sustaining either, as it is destroyed by any Pareto-efficient proposal.

[19] As usual, player $-i$ is allowed to use a mix strategy to minmax i.

frontier—consisting of vectors which are not strictly Pareto dominated—is a piecewise linear curve joining P_1 and P_2.

For any payoff vector X of $V(\delta)$, let $\pi_i(X)$ denote the i^{th} component of X. Thus, in the statement to follow, $\underline{v}_1 = \pi_1(P_2)$ denotes the first component of P_2 and $\underline{v}_2 = \pi_2(P_1)$ denotes the second component of P_1.

THEOREM 1

- *(Sufficiency) If*

$$\pi_i(A) > v_i \qquad \textit{for } i \in \{1,2\} \tag{1}$$

or $A = P_1 = P_2$, then $A \in V(q)$ for all $q \geq 0$.

- *(Necessity) If $A \in V(q)$, then*

$$\pi_i(A) \geq \underline{v}_i + \frac{q}{2+q}(\pi_i(P_{-i}) - \underline{v}_i) \tag{2}$$

for $i \in \{1,2\}$. If A is sustainable, inequalities in (1) *must hold for both players as weak inequalities.*

Condition (1) thus fully characterizes (up to its boundary) the set of sustainable payoffs. The sufficient and necessary conditions are respectively derived in Appendices A.1 and A.2.[20]

Figure 1 illustrates Theorem 1 for a fixed q: the green region represents the payoff vectors known to be sustainable and the orange region represents the additional vectors which may be sustainable. When $q = 0$ (no renegotiation), the orange region extends all the way to the minmax vector $\underline{v}$ and we recover the Folk Theorem. As renegotiation frictions become arbitrarily small ($q \to +\infty$), the orange region disappears as necessary and sufficient conditions become identical (up to their boundary).

One consequence of Theorem 1 is that $V(q)$ is nonempty for all $q \geq 0$ and so is the set of sustainable payoffs. In particular, our concept of self-sustainability provides a well-defined counterpoint to the standard Folk Theorem when renegotiation is introduced to repeated games, allowing us to compare the impact of renegotiation across different strategic situations of the stage game, from perfectly aligned interests to extreme misalignments, and to establish for a large class of games the possibility

[20] Appendix A.1 focuses on the case $P_1 \neq P_2$. If $P_1 = P_2$, players have perfectly aligned interests as they both want to implement P_1 and the construction is trivial. When $P_1 = P_2$, the necessary condition selects this vector as the unique outcome as renegotiation frictions become negligible. If the weak Pareto frontier lies strictly above the minmax values, and consists of a segment giving a constant payoff to one of the players—a degenerate case—any payoff on the frontier is sustainable.

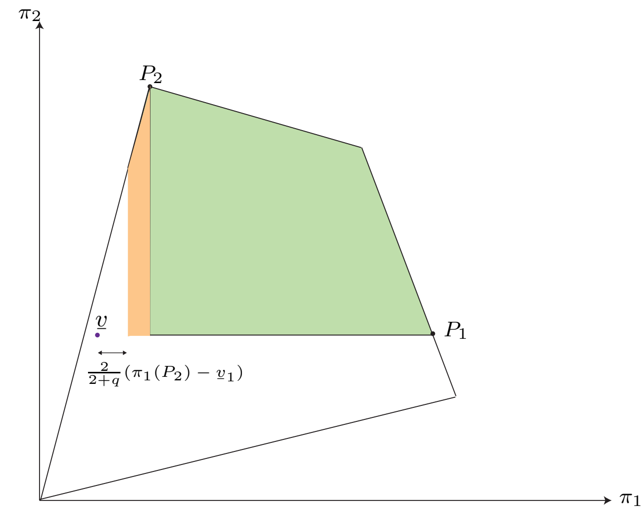

Figure 1: Necessary and sufficient conditions for fixed q

of sustaining inefficient norms even when δ is arbitrarily close to 1, and players can frictionlessly and credibly propose and agree on Pareto improving proposals.[21]

3.1 Relation between player alignment and sustainable outcomes

Figure 2 represents the set of sustainable payoffs for degrees of player alignment. In configuration (a), renegotiation constrains the set of implementable payoffs because the deterrence points P_1 and P_2 are too close to each other relative to the vector of minmax payoffs. Configuration (b) represents a perfectly cooperative game. The only sustainable outcome is the Pareto efficient vector. In configuration (c), the punishment/reward vectors used to deter off-path proposals are sufficiently far apart and the Folk Theorem holds despite the presence of frictionless renegotiation.

As the figure illustrates, the impact of renegotiation hinges on the alignment structure of the stage game. As the game becomes less cooperative (moving from (b) to (a) to (c) on the figure), there is more scope for disagreement among the players, which can be used to implement a larger set of feasible payoffs. Strategic renegotiation thus does not destroy the implementability of Pareto-efficient payoffs, but does not prevent Pareto-inefficient ones either, and the severity of the inefficiency which may be sustained increases as players' interests become more divergent.

[21] It should be noted that for fixed ε, there need not exist any self-sustaining norm, just as strongly renegotiation-proof equilibria (Farrell and Maskin (1989)) and externally consistent norms (Bernheim and Ray (1989)) may fail to exist for fixed discount factors. Indeed, we have constructed a family of counter-examples for some fixed $\varepsilon > 0$ and all values of $q > 0$.

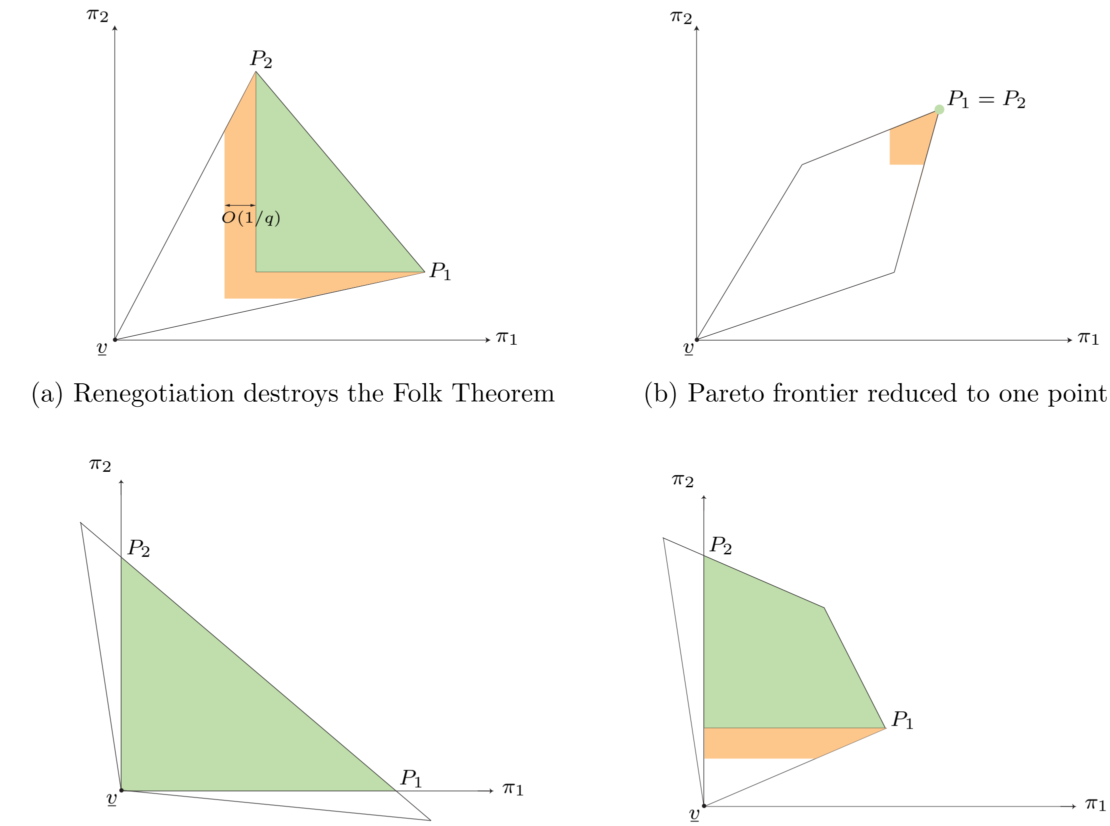

(a) Renegotiation destroys the Folk Theorem

(b) Pareto frontier reduced to one point

(c) Folk Theorem with extreme deterrence points

(d) Asymmetric case

Figure 2: Necessary and sufficient conditions for various configurations

3.2 Comparative statics: bargaining frictions and discounting

In standard repeated games with public randomization, it is well known that the set of implementable payoffs gets larger as δ limits to one. This property does not hold with renegotiation. For example, suppose that the stage game has an inefficient Nash equilibrium that violates the necessary conditions obtained by Theorem 1 for $q = \frac{1}{2}$. For small ε, Theorem 1 implies that this Nash equilibrium payoff, and an open neighborhood around it, is not sustainable. However as ε goes to 1, there is a norm under which players follow this Nash equilibrium in the first period (before possibly renegotiating to a Pareto superior continuation). Since the current-period weight is arbitrarily close to 1, players' payoffs are arbitrarily close to the inefficient Nash equilibrium's payoffs, which was impossible with a small enough value of ε.

Although discount-factor monotonicity is violated in the presence of renegotiation, a different kind of monotonicity arises here, with respect to negotiation frictions: the more opportunities players have to renegotiate their norm, the smaller the set of sustainable payoffs. This result holds at all discount factor levels and is proved in Appendix H.

Proposition 1 *For any fixed $\varepsilon \in (0,1)$, the set of sustainable payoffs is decreasing in q.*

3.3 Relation to the existing literature

When renegotiation is viewed as a strategic interaction, renegotiation-proof equilibria may contain Pareto-ranked continuations. This happens when Pareto-improving proposals are dissuaded by punishing the proposer and rewarding the rejector beyond the proposal. This idea also underlies the results of Santos (2000) who considers players bargaining over which equilibrium to play in a one-shot game, as well as Miller and Watson's (2013) Theorem 1, which shows that renegotiation has no restrictive power when it must only obey their "Internal Agreement Consistency" Axiom. That theorem and ours differ in two important ways. First, their argument requires unbounded transfers: to punish a proposer, say player 1, one requires him to make a very high transfer to 2 in the next period. If the weight of a single period is ε, the transfer must be of order $\frac{1}{\varepsilon}$, hence the necessity of unbounded transfers as ε goes to zero. These large transfers permit 1's continuation value to jump immediately from some punishment payoff v_1^0 to a higher continuation value v_1, which is easy to implement. Second, the transfer stage takes place, in each period, before the action stage (and, in particular, is distinct from it). If 1 deviates by making a lower transfer than prescribed, it suffices to have him minmaxed by the other player and reset the continuation value to v_1^0 for the next period in order to punish this deviation.

When stage-game payoffs are bounded, as in our setting, the continuation value of a player cannot jump by an ε-independent amount. The equilibrium construction must thus keep track of continuation values and make sure that these continuation values are implementable at all stages and following all deviations. In the absence of a separate transfer stage, moreover, if player 1 deviates in action when implementing v_1^0, his continuation value must fall below v_1^0. Implementing this lower value may be difficult or even impossible. In fact, it is this impossibility which creates new restrictions on the set of sustainable payoffs and destroys the Folk Theorem obtained in Miller and Watson's Theorem 1.

Both Santos (2000) and Miller and Watson (2013) consider a further restriction, which is that the continuation of the game, in case of a disagreement, be independent of the identity of the proposer and of the nature of the proposals.[22] This restriction guarantees a higher level of efficiency. The

[22]Similar ideas appear in Farrell (1987), Rabin (1994), and Arvan, Cabral, Santos (1999) for the case of simultaneous announcements.

consequences for our model of such a refinement are studied in Section L.

4 Interpretation and applications: a novel mechanism for miscoordination and inertia

The statement of Theorem 1 implies even when i) the game has complete information, ii) the discount factor δ is arbitrarily close to 1, iii) players can exchange messages at arbitrarily high frequency, the meaning of which consists of unambiguous agreement proposals, and iv) accepted agreements are binding.

Inefficient payoffs are sustained by rules which discourage proposals, including Pareto-improving ones, and which may be viewed as part of a social norm among the players. Although such a social norm seems undesirable from the perspective of the players, it may be interpreted in a broader context, in which the designer of the norm is not an active player of the game and benefits from the inefficiency that arises from the perspective of the players.

In particular, the agents who are explicitly modeled—only two so far, although Section 6 extends the analysis to an arbitrary number of players—may be part of a larger society or organization who exert externalities on other economic, unmodeled agents.

Potential applications include bidders in an auction, firms in a cartel, members of a radical organization, or simply citizens which the social planner wish to control, as in the case of a dictatorship. In these applications, high payoffs for these players mean that they are colluding, polluting, shirking, or, more generally, adversely affecting individuals who enter the social planner's objective.

Viewed from this perspective, the rules which enforce a Pareto-inefficient norm have the virtue of being self-sustaining: the social planner does not need to intervene once the game has started. There is no need for external monitoring or punishment.

Consider, for instance, a regulator wishing to prevent collusive pricing in an oligopolistic market. If the firms can be given self-sustaining rules that prevent collusion, such a design is of course cheaper for the regulator than explicitly monitoring the firms and administrating the punishments. Likewise, the manager of administrative office facing high costs of monitoring his employees may wish to create a social norm between them which implements high effort and under which an employee's proposal to shirk is rebuked by other employees and thwarted without requiring the manager's intervention. The designer's role consists in setting the rules at the beginning of the game, specifying how players should interpret deviations in actions and proposals. Once this common understanding is reached, the designer completely withdraws from the game: the players enforce the rules themselves by punishing one another if one of them ever deviates from these rules.

Of course, proposal-deterring norms do not have to be designed by anyone: players may simply be trapped in a norm with this feature—perhaps the remain of an unmodeled evolution before which such a norm made sense. An example may be "acting tough" and discouraging any suggestion to "soften up" even when doing so would in fact lead to a Pareto improvement.

Suppose that the designer of the norm can also make new actions available to the players (such as snitching on one another, as in the prisoner's dilemma). Introducing actions which increase the misalignment between players, as discussed in Section 3.1, increases the maximal inefficiency of sustainable outcomes, and thus potentially also the (unmodeled) payoff of the norm designer. From a designer's perspective, *creating* actions which benefit only one player but not others facilitates the deterrence of collusive proposals.

We sketch two applications below, in which norms are Pareto inefficient from the players' perspective but beneficial to some social planner whose sole involvement in the game, possibly, is to design the norm governing players' interactions.

Cournot competition.

Consider two symmetric firms which, under Cournot competition, produce together more than the monopolistic output. These firms could achieve a higher profit by each producing half of the monopolistic output. However, proposals to move away from the current equilibrium may be subject to a norm treating any such proposal as corrupt behavior. The firm on the receiving end of such a proposal could reject it, triggering a continuation in which, say, the rejector produces the Stackelberg leader's output in each period and the proposer produces the Stackelberg follower's output. These outputs give the proposer a lower payoff than the competitive equilibrium and his competitor a higher payoff than the half of the monopoly's profit.[23]

Political inertia and dictatorship.

Consider an authoritarian regime facing the risk of a revolution. In this regime, citizens may be exploited through high taxes, expropriation, and other channels. Faced with this situation, various citizen factions may attempt to persuade others to start a revolution (an off-path proposal). If the proposal is accepted, the authoritarian regime falls which (ideally) increases all citizens' well-being.

[23]The punishment for the proposer, i.e., the Stackelberg equilibrium, is inefficient. However, it suffices to incentivize a rejection to the proposal and thus deter a Pareto-improving proposal.

The regime may impose a norm that thwarts this threat by rewarding anyone who reveals the plot and punishing its instigator. Importantly, all rewards and punishments are *administered by the citizens, without the dictator getting involved or even monitor them.*[24] This provides a novel, completely endogenous explanation for the stability of dictatorships even when citizens can credibly coordinate to overthrow the regime. It exposes the limits of attempts to coordinate when the norm in place anticipates such attempts.

The dynamic nature of social norms and the importance of neologisms

Our analysis emphasizes the dynamic nature of social norms, particularly with regard to how "innovative" proposals are perceived. For instance, starting from a Pareto-inefficient norm, a first proposal to increase cooperation and increase both players' payoffs may perceived as acceptable and implemented. To be sustained, however, this cooperation may require the threat of punishments during which the kind of cooperation originally proposed is no longer acceptable.

In other applications, a dynamic norm may capture each player's endogenous "status" determining the actions and proposals that he is allowed to pursue. If a player proposes to disrupt the current norm (say, by implementing a higher cooperation, or a revolution), he loses his status when the proposal is rejected. An infamous example concerns prisoner camps in which a prisoner is assigned the administrative authority over other prisoners, but loses it if this authority is used to rebel.

5 Equivalent notions of stability

This section provides "open" and "closed" set-theoretic formulations of self-sustaining norms, and then shows that our necessary and sufficient conditions are unchanged if only a subset of "credible" proposals is taken seriously.

5.1 Stable Conventions

The first one defines a set-theoretic notion of norms, which facilitates the comparison of our concept with existing notions of renegotiation-proofness (e.g., Farrell and Maskin (1989) and Bernheim and Ray (1989)).

Definition 2 *A subset $\mathcal{C} \subset \mathcal{S}$ of norms is a* **convention** *if for any $s \in \mathcal{C}$ as long as no off-path*

[24] While the application obviously involves more than two players, the gist of the many-player analysis is identical to the two-player one, as shown in the many-player extension of Section 6.

proposal was accepted, the continuation of s belongs to the convention $\mathcal{C}$.[25]

The definition implies that if players start with a norm in a convention, then all on-path proposals (whether they are accepted or rejected), as well as rejected off-path proposals, have their continuations in the convention. In particular, deviations in actions are punished within the convention, as long as no off-path proposal to leave the convention has been accepted. One may informally view $\mathcal{C}$ as a "social norm:" it describes the set of continuations which players perceive as consistent with "business as usual." A convention can be abandoned only when some player makes an off-path proposal outside of it that is accepted by the other player. The following notion of stability requires that such proposals be taken seriously by the players.

Definition 3 *A convention $\mathcal{C}$ is* **stable** *if, in any period starting with a norm in $\mathcal{C}$, the properties of Message Richness and Binding Acceptance from Definition 1 are satisfied.*

Since all on-path continuations of norms in $\mathcal{C}$ must all belong to $\mathcal{C}$—by definition of a convention—stability implies that any Pareto-improving proposal lying outside the convention is rejected with probability 1; for if it were accepted, stability would require that the proposal be implemented. Stability thus requires that no player ever has an incentive to make proposals outside of the convention—hence the terminology. Intuitively, stability is achieved by rewarding a player on the receiving end of a deviating proposal whenever he rejects it. Crucially, however, this continuation, which rewards the rejector and deters the proposer, must lie *within the convention.*

As anticipated, convention stability is equivalent to self-sustainability in the following sense.

Proposition 2 *An norm is self-sustaining if and only if it is part of a stable convention.*

The proof of this equivalence is straightforward: First, any norm belonging to a stable convention must be self-sustaining, since all continuations of on-path proposals and rejected off-path proposals lie in the convention, and thus subject to the stability condition. For the reverse direction, take any self-sustaining norm and consider the set consisting of this norm together with all of its continuations at the beginning of periods following histories for which no off-path proposals has been accepted. This set forms a convention, by construction, which is stable, by self-sustainability of the norm.

Closed vs. open conventions

[25] Continuations of s are always defined at the beginning of the corresponding period.

The conventions defined above are *open* in the sense that they allow players to depart from the convention when an off-equilibrium proposal is accepted. This possibility is absent from earlier studies of renegotiation-proof equilibrium. However, these perspectives can be reconciled: we show that convention stability can be recast in terms of a purely set-theoretic definition.

Definition 4 *A subset $\mathcal{C}$ of $\mathcal{S}$ is a* **closed convention** *if for any $s \in \mathcal{C}$, any continuation of s belongs to $\mathcal{C}$.*

The only difference with Definition 2 is that continuations belong to the convention even when off-path proposals are accepted. To offset this change, our earlier definition of stability is translated into the language of set-theoretic analysis.

Definition 5 *A closed convention $\mathcal{C}$ is* **stable** *if it satisfies the following property: Consider any norm of $\mathcal{C}$ and history at which i gets a chance to make a proposal, and let $\hat{U}_i$ denote i's continuation payoff. Then, for any proposal with payoff vector U which gives i a payoff $\pi_i(U) > \hat{U}_i$, there is a payoff vector U' of $\mathcal{C}$ such that $\pi_{-i}(U') \geq \pi_{-i}(U)$ and $\pi_i(U') \leq \hat{U}_i$.*

Theorem 2

1. For any closed convention $\mathcal{C}^c$, there exists an open convention $\mathcal{C}^o$ which has the same payoff set, and vice versa.

2. For any stable closed convention $\mathcal{C}^c$, there exists a stable open convention $\mathcal{C}^o$ which has the same payoff set, and vice versa.

5.2 Credible proposals

Stability requires that players implement any accepted proposal. When players are used to a convention $\mathcal{C}$, one may wonder why players should take all proposals seriously, particularly when these proposals lie outside of the convention. It turns out that Theorem 1's necessary and sufficient conditions are identical if one restricts proposals to a much smaller subset of "credible" proposals.

Definition 6 *Given a convention $\mathcal{C}$, a norm is $\mathcal{C}$-credible (or just "credible", when there is no confusion) if any off-equilibrium play (action, proposal, or acceptance decision) triggers a continuation that belongs to $\mathcal{C}$ (for the appropriate stage within the period). A continuation payoff is credible if it is implemented by a credible norm.*

Starting with a norm belonging to some convention $\mathcal{C}$, a credible proposal is such that any future deviation triggers a reversal to the convention. For example, if a convention includes a "punishment" norm that gives low utility to both players, the convention can support many credible norms by imposing that any deviation trigger the punishment norm. Since the players may deviate at different stages of any period, after a deviation they will play the next stage according to the convention.[26] In addition, any subsequent deviation (namely, accepting an off-path proposal) may trigger a norm which does not belong to the convention.

Definition 7 *A convention $\mathcal{C}$ is credibly stable if:*

1. **Message $\mathcal{C}$-Richness** *Each $m \in \mathcal{M}$ is associated with a norm of $\mathcal{S}$, and each $\mathcal{C}$-credible payoff is implemented by an $\mathcal{C}$-credible norm associated with some message;*

2. **$\mathcal{C}$-Binding Acceptance** *If a message is accepted, whose associated norm is $\mathcal{C}$-credible, the associated norm is implemented.*

Definition 7 is clearly more permissive than Definition 3, because it imposes the refinement over a smaller set of proposals. However, we have the following result.

Theorem 3 *The set of points sustained by credibly stable conventions obeys the necessary and sufficient conditions of Theorem 1.*

The proof is straightforward: first, any stable convention is credibly stable since the latter must sustain fewer proposal challenges than the former. Our construction for the sufficiency condition thus still applies. Second, the proposals used in Appendix A.2 to derive the necessary conditions of Theorem 1 are credible, as shown in this appendix. The necessary conditions are thus identical for stable and credibly stable conventions.

6 Arbitrary number of players

The analysis so far has focused on two players, a common restriction to study renegotiation in repeated games.[27] Extending the analysis to more players raises new conceptual issues. Can

[26]For example, if a player deviates during the action stage, the players will then engage in renegotiation under the rules prescribed by the convention. If a player deviates during the renegotiation stage by sending the wrong message or making the off-path acceptance choice, then in the next period the players will choose their actions according to the convention.

[27]E.g., Farrell and Maskin (1989), Benoît and Krishna (1993), and Santos (2000). Abreu, Pearce, and Stacchetti (1993) focus instead on symmetric equilibria.

proposals be targeted toward a subset of individuals? What happens if only a subset of the players accepts the proposal?

This section explores some of these issues, allowing for an arbitrary number, n, of players. After a player has made a proposal, other players vote on accepting the proposal. We assume that the vote is simultaneous and show in an extension that sequential voting does not alter our conclusions (Appendix J.2).

The setting build on the two-player case is modified as follows. At the proposal stage, each player i has the same probability $\frac{p}{n}$ $(p < 1)$ of being chosen to send a message. The renegotiation friction parameter q is still defined by $p = q\varepsilon$. This player may choose to conceal his opportunity to send a message. If i sends a message, other players vote on whether to accept it, resulting in a vector of acceptance votes $D_{-n} \in \{0, 1\}^{n-1}$.[28]

With multiple players voting on a proposal, we consider the *supermajority* rule: a proposal is accepted if at least L players support it, with $L \in \{\lfloor N/2 \rfloor, \ldots, N-1\}$.[29] Each fixed value of L defines a concept of self-sustaining norm refinement, as in the two-player case. Let $\mathcal{H}^+$ denote the set of all finite public histories ending after an action stage such that no off-path proposal has been accepted by the supermajority.

Definition 8 *An norm is* self-sustaining *if the following holds for any history $h \in \mathcal{H}^+$:*

1. **Message Richness** *Each message $m \in \mathcal{M}$ is assigned a norm of $\mathcal{S}$. The payoff vectors associated to these norms cover the set $V(\delta)$;*
2. **Binding Acceptance** *If a message is accepted by at least L voters, the norm assigned to it is played from next period onward.*

As with with most of the literature on voting, we assume that if a proposal fails the vote, the continuation is independent of the exact number, or identity, of the voters who voted to reject it.[30]

Definition 9 *A self-sustaining norm s is* **simple** *if, for any history $h \in \mathcal{H}^+$, when player i*

[28] As in the two-player case, if no message is sent the identity of a sender is arbitrarily chosen and the empty message is assumed to be rejected by everyone else.

[29] The lower bound $N/2$ is natural to interpret the voting as a supermajority rule, but not necessary for the analysis. The upper bound $N-1$ corresponds to the unanimity rule, keeping in mind that the proposer is not voting over his own proposal.

[30] The case in which continuations can depend arbitrarily on the voting profile in considered in Appendix J.1.

makes a proposal m_i, there are two continuations, depending on whether m_i passes or fails the supermajority vote.

In the analysis that follows, we focus to fix ideas on the unanimity rule ($L = N - 1$). This rule is easier to interpret (since no player is "forced" to espouse a new norm that he has not chosen), but the analysis of other supermajority rule is qualitatively the same. As usual with voting games, we eliminate weakly dominated strategies.

ASSUMPTION 1 *A player votes in favor of the proposal if it gives him a strictly higher payoff than its continuation payoff in case of a rejection.*

We also assume that the individually-rational payoff set has full-dimension, which guarantees that the Folk Theorem holds for the underlying repeated game (Fudenberg and Maskin (1986)).

The key to characterizing self-sustaining norms is to determine each player i's worst possible punishment if he makes an unprescribed proposal. Suppose that i makes a proposal with corresponding payoff vector C, and let $\mathcal{V}$ denote the set of payoff vectors across all continuations of our candidate self-sustaining norm, s, following histories at which no off-path proposal has been accepted by the supermajority. (All payoffs in $\mathcal{V}$ are estimated at the *beginning* of a period.) If s is self-sustaining, C will be implemented if all players accept i's proposal. If anyone rejects the proposal, norm simplicity implies that there is a single payoff vector *in* $\mathcal{V}$, $D(C)$, which will be realized. If $D(C)$ gives $\pi_j(C)$ or more to at least one player $j \neq i$, this player will refuse the implementation of C, and the norm implementing $D(C)$ will be played.

Following any proposal with payoff C by player i, the worst punishment in $\mathcal{V}$ for player i minimizes i's utility over the set:[31]

$$\mathcal{D}(C, \mathcal{V}) = \{D(C) \in \mathcal{V} : \exists j \neq i : \pi_j(D(C)) \geq \pi_j(C)\}.$$

Let $\underline{\pi_i}(C, \mathcal{V})$ denote i's utility under this worst punishment.

Viewing $\underline{\pi_i}(C, \mathcal{V})$ as a function of C, one can then find the proposal with a continuation $C(\mathcal{V})$ which maximizes i's payoff at the worst punishment: $C(\mathcal{V}) = \arg\max_C\{\underline{\pi_i}(C, \mathcal{V})\}$, and the corresponding payoff, $\underline{\pi_i}(\mathcal{V})$, for i. The payoff $C(\mathcal{V})$ may be viewed as follows. The most efficient way of preventing player i from making a non-prescribed proposal is by implementing his worst punishment. Anticipating this, if player i deviates from his prescription, he may as well choose the optimal proposal, which gives the payoff $C(\mathcal{V})$.

[31] For the existence of a *worst* punishment, the set $\mathcal{V}$ needs to be closed. Our construction will satisfy this condition.

These observations lead to the following sequential construction. We start from the set $\mathcal{F}$ of strictly individually-rational payoffs in the stage game, i.e., the payoffs which would be implementable without renegotiation as δ goes to 1. We then consider the minimal payoffs $\underline{\pi_i}(\mathcal{F})$, $i \in \{1, ..., n\}$ that each player i could guarantee himself when given a chance to make a proposal if all sustainable payoffs belong to $\mathcal{F}$. We will build *two* decreasing sequences of sets, starting from $\mathcal{F}$, which will generate separate necessary and sufficient conditions for a payoff to be self-sustaining.

To derive sufficient conditions, the k^{th} set in the sequence, $\mathcal{F}_S^k$, is reduced by removing all the payoffs below $\underline{\pi_i}(\mathcal{F}_S^k)$, to form the $k+1$-th set in the sequence, starting with $\mathcal{F}_S^0 = \mathcal{F}$. We will show that this process converges to a stable set which defines sufficient conditions.

To derive necessary conditions, the k^{th} set in the sequence, $\mathcal{F}_N^k$, is constructed inductively as follows. Let $\pi_{min,i}(\mathcal{F}_N^k)$ denote the lowest expected payoff for player i at the beginning of a period, among all payoff vectors in $\mathcal{F}_N^k$. This value is lower than the continuation payoff $\underline{\pi_i}(\mathcal{F}_N^k)$ that i can guarantee himself when he gets a chance to make a proposal. We have

$$\pi_{min,i}(\mathcal{F}_N^k) \geq \varepsilon \underline{v_i} + (1-\varepsilon)\left[\frac{q\varepsilon}{n}\underline{\pi_i}(\mathcal{F}_N^k) + (1-\frac{q\varepsilon}{n})\pi_{min,i}(\mathcal{F}_N^k)\right]$$

Indeed, as in the two-player case, i gets at least $\underline{v_i}$ as his current payoff, and can guarantee himself $\underline{\pi_i}(\mathcal{F}_N^k)$ if he has a chance to make a proposal. As ε goes to 0, one can express the value $\pi_{min,i}(\mathcal{F}_N^k)$ as:

$$\pi_{min,i}(\mathcal{F}_N^k) \geq \frac{n\underline{v_i} + q\underline{\pi_i}(\mathcal{F}_N^k)}{n+q}. \tag{3}$$

At each step the set $\mathcal{F}_N^k$ is being reduced by removing the payoffs below (3). Iterations of this procedure converge to a steady set, as we show in the Appendix.

Proposition 3 *Both procedures converge to steady sets.*

We denote the limiting sets by $\mathcal{V}_S$ and $\mathcal{V}_N$. These sets are both positive orthants, whose vertices give lower bounds on players' payoffs (calculated at the beginning of period) under both procedures, and are denoted $\pi_{min,i}(\mathcal{V}_S)$ and $\pi_{min,i}(\mathcal{V}_N)$, for any player i. By construction, expression (3) holds as an equality for $\mathcal{V}_N$:

$$\pi_{min,i}(\mathcal{V}_N) = \frac{n\underline{v_i} + q\underline{\pi_i}(\mathcal{V}_N)}{n+q} \tag{4}$$

Similarly, we have $\pi_{min,i}(\mathcal{V}_S) = \underline{\pi_i}(\mathcal{V}_S)$.

We can now state the main result of this section. Let $\mathcal{P}$ denote the Pareto frontier of the feasible payoffs in the stage game and, for each i, P_{-i} denote any individually-rational payoff vector of $\mathcal{P}$ which minimizes i's payoff. Also let $\mathcal{R}$ denote the open positive orthant whose vertex is the vector $(\pi_i(P_{-i}))_{i=1}^n$. In the two-player case, this set characterized the sufficient conditions for sustainable payoffs. With $n > 2$ players, we show that $\mathcal{R}$ still consists of sustainable payoff vectors, though it might not include all of them. The theorem is formulated for the case where Pareto frontier supports for each player a non-zero range of payoffs.[32]

Theorem 4 *Any sustainable payoff lies in $\mathcal{V}_N$, and generically any payoff in the interior of $\mathcal{V}_S$ is sustainable. Moreover, any payoff in the interior of $\mathcal{R}$ is sustainable.*

As in the two-player case, the sets defined by the necessary and sufficient conditions converge to each other as renegotiation frictions vanish (see Appendix C.3). Moreover, by construction, the two sets shrink if the number L that determines the supermajority rule, decreases.

Proposition 4 *The sets $\mathcal{V}_S$ and $\mathcal{V}_N$ converge to each other as q goes to infinity. The sets shrink as L decreases.*

When players respond sequentially to a proposal, the same result obtains.[33]

7 Conclusion

This paper provides a model of self-enforcing agreements in which proposals to overturn these agreements are explicitly considered as part of a larger game. Self-sustaining norms are equilibria of this enlarged game, i.e., self-enforcing agreements which concern not just which actions to take but also how to react to proposals to change the agreements. We characterize the set of self-sustaining norm payoffs as the discount factor goes to 1.

One virtue of our model is its simplicity, which is reflected in three aspects: i) the protocol of negotiation, which consists of a one-shot proposal/acceptance stage to the stage game, ii) the concept of self-sustaining norms, which is a single, straightforward equilibrium refinement, iii) the characterization of the set of sustainable payoffs, which are easily described graphically.

Another virtue is its robustness: the results are unchanged if the negotiation protocol is expanded, or if the refinement is restricted to a set of credible proposals. No assumption is imposed on the

[32]When the feasible set has a unique Pareto-efficient point, this point is the only sustainable payoff. When at least two players have multiple payoffs on the Pareto frontier, the set of sustainable payoffs is always a non-empty full-dimensional orthant. When all points of the Pareto frontier give the same payoff to all but one player, any payoff on the Pareto frontier is sustainable.

[33]See Appendix J.2. Sequential and simultaneous voting are in fact equivalent under the assumption, standard in the voting literature, that a player votes for the proposal if it gives him higher expected payoff than rejecting it.

nature of the stage game beyond the standard full dimensionality of the feasible set, when there are three or more agents.

Conceptually, our notion of self-sustaining norms formalizes the idea that the rules governing our interactions are subject to change but also subject to deeper norms that govern how such changes are perceived. This seems to capture the concept of "cultural beliefs" studied by Greif (1994) and provide a mechanism for the persistence of inefficiencies even when agents are free to communicate.

Beyond these results, some important issues remain to be explored. In particular, what happens if a player can make a proposal to a subset of players? How would such a proposal, if accepted, affect the strategies used by the players excluded from the proposal? Exclusive negotiations of this kind are common in economics, when agents are divided into relatively homogeneous groups within which negotiation is easier or when they are engaged in specific relationships like those arising in supplier chains. They may also arise in community enforcement models, in which matching parties may engage in local renegotiation to alleviate punishments (Ali, Miller, and Yang (2016)). Understanding how strategic renegotiation shapes equilibrium outcomes in environments with segmented groups seems a particularly interesting direction for future work.

References

Abreu, D., Pearce, D. (1991) "A perspective on renegotiation in repeated games," in R. Selten (ed.), *Game Equilibrium Models*, Springer Berlin Heidelberg.

Abreu, D., Pearce, D., and E. Stacchetti (1993) "Renegotiation and Symmetry in Repeated Games," *Journal of Economic Theory*, Vol. 60, pp. 217–240.

Ali, N., Miller, D., and D. Yang (2016) "Renegotiation-Proof Multilateral Enforcement," *Working Paper*, Pennsylvania State University and University of Michigan.

Arvan, L., Cabral, L., and V. Santos (1999) "Meaningful cheap talk must improve equilibrium payoffs," *Mathematical Social Sciences*, Vol. 37, pp. 97–106.

Asheim, G. (1991) "Extending Renegotiation-Proofness to Infinite Horizon Games," *Games and Economic Behavior*, Vol. 3, pp. 278–294.

Atkeson, A. (1991) "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica*, Vol. 59, pp. 1069–89.

Benoît, J.-P., Krishna, V. (1993) "Renegotiation in Finitely Repeated Games," *Econometrica*, Vol. 61, pp. 303–323.

Bergin, J., MacLeod, B. (1993) "Efficiency and Renegotiation in Repeated Games," *Journal of Economic Theory*, Vol. 61, pp. 42–73.

Bernheim, B.D., Ray, D. (1989) "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 295–326.

Bernheim, B.D, Peleg, B., Whinston, M. (1987) "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, Vol. 42, pp. 1-12.

Blume, A. (1994) "Intraplay Communication in Repeated Games," *Games and Economic Behavior*, Vol. 6, pp. 181–211.

Brennan, J.R., Watson, J. (2013) "The Renegotiation-proofness Principle and Costly Renegotiation," *Games*, Vol. 4, pp. 347–366.

Bulow, J., Rogoff, K. (1989) "A Constant Recontracting Model of Sovereign Debt," *Journal of Political Economy*, Vol. 97, pp. 155–178.

Bulow, J., Rogoff, K. (1991) "Sovereign Debt Repurchases: No Cure for Overhang," *Quarterly Journal of Economics*, Vol. 106, pp. 1219–35.

Caillaud, B. and Jehiel, P. (1998) "Collusion in Auctions with Externalities," *Rand Journal of Economics*, Vol. 29, pp. 680–702.

Cave, J. (1987) "Long-Term Competition in a Dynamic Game: The Cold Fish War," *Rand Journal of Economics*, Vol. 18, pp. 596–610.

Che, Y.-K., Kim, J. (2006) "Robustly Collusion-Proof Implementation," *Econometrica*, Vol. 74, pp. 1063–1107.

Crémer, J. (1996) "Manipulations by Coalitions Under Asymmetric Information: The Case of Groves Mechanisms," *Games and Economic Behavior*, Vol. 13, pp. 39–73.

DeMarzo, P. (1988) "Coalitions and Sustainable Social Norms in Repeated Games," *mimeo*, Stanford University.

Farrell, J. (1983) "Credible Repeated Game Equilibrium," *Unpublished Manuscript.*

Farrell, J. (1987) "Cheap Talk, Communication, and Entry." *Rand Journal of Economics*, Vol. 18, pp. 34–39.

Farrell, J., Maskin, E. (1989) "Renegotiation in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 327–360.

Fershtman, C., Nitzan, S. (1991) "Dynamic Voluntary Provision of Public Goods," *European Economic Review*, Vol. 35, pp. 1057–67.

Fudenberg, D., Maskin, E. (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, Vol. 54, pp. 533–554.

Greif, A. (1994) "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection onCollectivist and Individualist Societies," *Journal of Political Economy*, Vol. 102, pp. 912–950.

Levhari, D., Mirman, L.J. (1980) " The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution," *Bell Journal of Economics*, Vol. 11, pp. 322–334.

Levin, J. (2003) "Relational Incentive Contracts," *American Economic Review*, Vol. 93, pp. 835–857.

MacLeod, B., Malcomson, J. (1989) "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, Vol. 57, pp. 447–480.

Mendelson, E (1997) *Introduction to Mathematical Logic*, 4th ed. Chapman & Hall, London.

McAfee, R.P., McMillan, J. (1992) "Bidding Rings," *American Economic Review*, Vol. 82, pp. 579–599.

Miller, D., Watson, J. (2013) "A Theory of Disagreement in Repeated Games with Bargaining," *Econometrica*, Vol. 81, pp. 2303–2350.

Pearce, D. (1987) "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," *Cowles Foundation Discussion Paper*, No. 855.

Rabin, M. (1994) "A Model of Pre-Game Communication," *Journal of Economic Theory*, Vol. 63, pp. 370-391.

Ray, D. (1994) "Internally Renegotiation-Proof Equilibrium Sets: Limit Behavior with Low Discounting," *Games and Economic Behavior*, Vol. 6, pp. 162–177.

Santos, V. (2000) "Alternating-Announcements Cheap Talk," *Journal of Economic Behavior & Organization*, Vol. 42, pp. 405–416.

A Proof of Theorem 1

A.1 Sufficient Conditions

We construct, for any payoff vector (hereafter, "point") A satisfying (1) and ε sufficiently small a norm $s \in \mathcal{S}$ which implements A and is self-sustaining at all frequencies $q \geq 0$. The construction is based on points A_1 and A_2 such that A_i gives i his worst possible payoff among all *self-sustaining* continuations of s (that is, continuations following histories at which no off-path proposal was accepted).[34] When i's continuation payoff is at an ε-independent distance above his payoff from A_i, it is easy to incentivize him to follow any prescribed action, since any deviation provides a maximal gain of order ε and can be punished by implementing A_i. One challenge is to choose A_i so that i is adequately incentivized near A_i. The second important points of the construction are D_1 and D_2, which serve to deter off-path proposals. These points are chosen to be Pareto efficient, and set so that any relevant off-path proposal by i may be deterred by having $-i$ reject the proposal and have D_i be implemented instead. D_i must therefore be chosen so that $-i$ is sufficiently rewarded, and i punished, for any proposal that i may entertain.

Preliminaries

Since the message space has the cardinality of the continuum, we can without loss of generality identify it with the set $V(\delta)$ of feasible payoff vectors, a full dimensional subset of $\mathbb{R}^2$.[35]

We interpret each message $X \in V(\delta)$ as a proposal to move to a continuation whose expected payoff is X. For any point X implemented by some self-sustaining continuation of the candidate norm s, let s^X denote the corresponding continuation.[36]

To distinguish players' expected payoffs at each stage of each period, we introduce the following notation. Given a subset $\mathcal{L}$ of norms, let $\mathcal{U}(\mathcal{L}) \subset \mathbb{R}^2$—or just $\mathcal{U}$, when there is no confusion—denote the set of expected payoffs for the players across all possible norms in $\mathcal{L}$, computed before public randomization. $\mathcal{V}$ is defined identically but computed after the realization of the randomization device z. $\mathcal{U}$ is thus included in the convex hull of $\mathcal{V}$. Finally, let $\mathcal{W}$ consist of continuations payoffs after actions and payoffs are observed and incurred in the current period, but before the proposal stage. Each element of $\mathcal{W}$ is a convex combination of three expected payoff vectors corresponding to the following events: player 1 gets to make a proposal, player 2 does, or no one does. Because elements of $\mathcal{W}$ define continuation payoffs excluding the current period, to make them

[34] Unless stated otherwise, points refer to expected payoffs at the beginning of the current period.

[35] If $\mathcal{M}$'s cardinality exceed the cardinality of the continuum, we assign the minmax payoff to all superfluous messages.

[36] There is only one continuation for each payoff X considered below, so s^X is well defined.

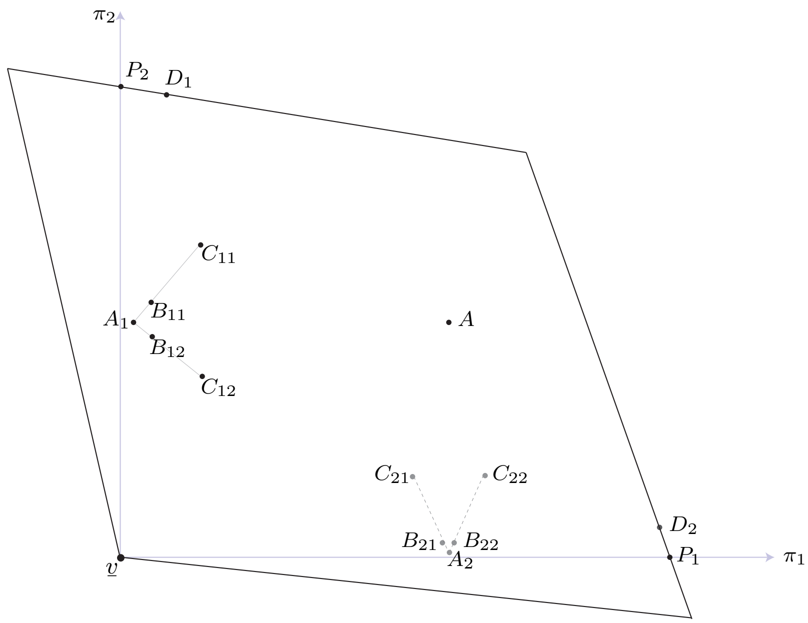

Figure 3: Construction of a self-sustaining norm

commensurate with payoffs in $\mathcal{U}$, we evaluate them at the next period (i.e., ignoring the discount factor between the two periods). With this convention, payoffs in $\mathcal{W}$ are convex combinations of elements of $\mathcal{U}$.

Elements of $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{W}$ are points of two-dimensional sets. Recall that for any payoff vector X, we let $\pi_i(X)$ denote the i^{th} component of X.

Proof

For each player i, there are two configurations to consider, depending on whether i's minmax payoff $\underline{v}_i$ is less than or equal to $\pi_i(P_{-i})$. We first consider the case in which both players are in the former configuration.

Case 1: $\underline{v}_1 = \pi_1(P_2)$ **and** $\underline{v}_2 = \pi_2(P_1)$ Consider any point A satisfying (1). For ε small enough, the points A_1 and A_2 with coordinates $(\pi_1, \pi_2)(A_1) = (\underline{v}_1 + \varepsilon^{\frac{1}{2}}, \pi_2(A))$ and $(\pi_1, \pi_2)(A_2) = (\pi_1(A), \underline{v}_2 + \varepsilon^{\frac{1}{2}})$ are individually rational and such that $\pi_1(A_1) < \pi_1(A)$ and $\pi_2(A_2) < \pi_2(A)$. The norm s^{A_1} implementing A_1 is constructed as follows (s^{A_2} has a similar construction):

1) Action stage: player 2 minmaxes player 1, possibly mixing between several actions $\{a_{2j}\}_j$. Player 1 best responds by a pure action $a_{1,minmax}$ achieving his minmax payoff.

1a) If no deviation in action is observed, the continuation payoff vector $B_{1j} \in \mathcal{W}$ is a function of 2's realized action, a_{2j}, and is chosen so that i) 2 is indifferent between all actions a_{2j} used to minmax

1, ii) 1's continuation payoff is independent of j (so the vectors $\{B_{1j}\}_j$ all lie on the same vertical line as shown on Figure 3), and iii) the promise-keeping condition is satisfied for both players. In particular,

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + (1-\varepsilon)\pi_1(B_{1j}) \tag{5}$$

for all indices j corresponding to some action a_{2j} in 2's minmaxing distribution. In particular, the points B_{1j} all lie within an ε-proportional distance of A_1.

1b) If 2 deviates in action (i.e., chooses an action outside of the mixture used to minmax 1), the continuation payoffs jump to the point A_2, mentioned above, which gives 2 her lowest possible payoff.[37] For small ε, this punishment suffices to incentivize 2 because any deviation gain is of order ε whereas $\pi_2(A_2)$ is arbitrarily close to 2's minmax payoff, causing 2 an ε-independent loss.

1c) If 1 deviates in action, disregard this. Such a deviation is suboptimal since 1 was prescribed to best respond to being minmaxed by 2.

2) Proposal stage: the norm $s^{B_{1j}}$ implementing B_{1j} is as follows: if either 2 gets a chance to make a proposal, or no player does, the play returns to s^{A_1}. 2 is prescribed to remain silent. If 1 gets a chance to make a proposal, he proposes a continuation $s^{C_{1j}}$ whose corresponding payoff vector C_{1j} lies on the line going through A_1 and B_{1j} and is chosen so as to satisfy the promise-keeping condition

$$\pi_1(B_{1j}) = \left(1 - \frac{p}{2}\right)\pi_1(A_1) + \frac{p}{2}\pi_1(C_{1j}) \tag{6}$$

Player 2 is prescribed to accept proposal $s^{C_{1j}}$. The points $\{C_{1j}\}_j$ give the same payoff to 1, independently of j. Their implementation is described in 3) below.

2a) If 1 proposes any continuation other than $s^{C_{1j}}$ that improves his payoff, he is punished by a continuation s^{D_1}—triggered if player 2 rejects 1's proposal—chosen such that i) $\pi_1(D_1) < \pi_1(C_{1j})$ and ii) 2 prefers $\pi_2(D_1)$ to her payoff under 1's proposal $s^{C_{1j}}$. Precisely, D_1 is defined as the point of the Pareto frontier that gives 1 a payoff of

$$\frac{\pi_1(A_1) + \pi_1(C_{1j})}{2} \tag{7}$$

As explained shortly, 1's payoff at C_{1j} is of order $\sqrt{\varepsilon}$ above what 1 gets at A_1 or B_{1j}. If 1 proposes a plan that makes him worse off than $s^{C_{1j}}$, 2 accepts it if only if improves her payoff. Of course, such a proposal never arises in equilibrium.

2b) If 2 deviates by making a proposal or rejecting 1's offer to move to C_{1j}, players jump to the continuation s^{D_2}, which punishes 2's deviation (in the former case, s^{D_2} is assigned as the

[37] More precisely, players start implementing the payoff B_{21}, which is the analogue of the point B_{11}, following the implementation of A_2.

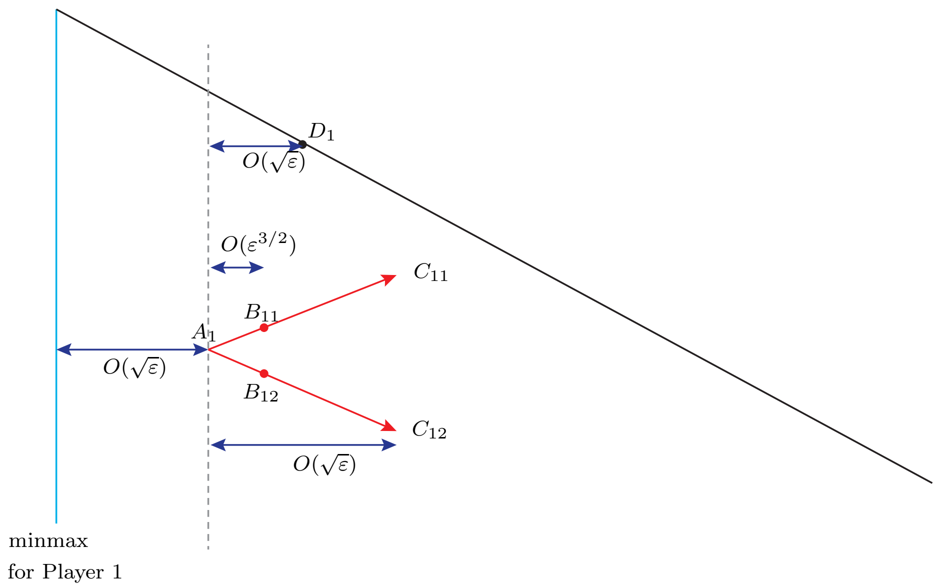

Figure 4: Construction details

continuation arising when 1 rejects 2's proposal). No proposal simultaneously gives 1 more than his payoff at s^{D_2} and 2 more than his payoff from s^{A_1}.

3) Next periods: the norm $s^{C_{1j}}$ is easily sustained because it gives 1 a payoff of order $\sqrt{\varepsilon}$ above what A_1 and B_{1j} give him. A deviation in action by 1 brings a gain of order ε and is punished by a drop of order $\sqrt{\varepsilon}$ in 1's continuation payoff, and is thus suboptimal, for ε small enough. $s^{C_{1j}}$ can be implemented by a deterministic sequence of actions keeping players' continuation payoffs within a distance $K\varepsilon$ from C_{1j}. The rules implementing this sequence are simple: play a deterministic action profile keeping continuation payoffs ε-close to C_{1j} and do not allow any proposal. If 1 deviates in actions, jump to one of the continuations $s^{B_{1j}}$; if he deviates in proposals, jump to s^{D_1} if 2 rejects this offer. A similar rule is applied for player 2, who has even more to lose from a deviation.

4) The point D_1 also gives 1 a payoff of order $\sqrt{\varepsilon}$ above A_1 and B_{1j}. s^{D_1} can therefore be implemented similarly to $s^{C_{1j}}$. Again, any proposal is ignored.

The construction is represented on Figure 3. The magnitudes of payoff differences between the points involved in the construction are indicated on Figure 4.

We verify the claim that all C_{1j}'s lie at a $\sqrt{\varepsilon}$-proportional distance to the right of A_1. From (5) and (6), we have

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + (1-\varepsilon)\pi_1(B_j) = \varepsilon \underline{v}_1 + (1-\varepsilon)\left[\left(1-\frac{q\varepsilon}{2}\right)\pi_1(A_1) + \frac{q\varepsilon}{2}\pi_1(C_{1j})\right]$$

Ignoring the terms of order ε^2 and higher, this implies that

$$\pi_1(A_1) = \varepsilon \underline{v}_1 + \left(1 - \left(1 + \frac{q}{2}\right)\varepsilon\right)\pi_1(A_1) + \frac{q\varepsilon}{2}\pi_1(C_{1j}).$$

Subtracting $\pi_1(A_1)$ from both sides and dividing by ε yields

$$\varepsilon^{\frac{1}{2}} = \pi_1(A_1) - \underline{v}_1 = \frac{q}{2}\left(\pi_1(C_{1j}) - \pi_1(A_1)\right), \tag{8}$$

which shows the claim.

The direction of each vector $\overrightarrow{A_1C_{1j}}$, which is also $\overrightarrow{A_1B_{1j}}$'s direction, depends only on 2's action, a_{2j}; it does not change when ε goes to 0. This shows that, for ε small enough, C_{1j} is a feasible payoff and $\pi_2(C_{1j})$ exceeds $\pi_2(A_2)$ by an ε-independent value.

As noted, the system of actions and proposals implementing s^{A_i}'s, $s^{B_{ij}}$'s and $s^{C_{ij}}$'s and s^{D_i}'s is incentive compatible in actions and in proposals. To conclude the construction, observe that A gives each player i a payoff higher than A_i, by an amount that is independent of ε. One may therefore implement A by a deterministic sequence of actions, chosen so that the continuation payoffs stay within a distance $K\varepsilon$ of A.[38] Deviations in actions are punished by moving to $s^{B_{11}}$ or $s^{B_{21}}$, depending on the deviator's identity. Deviations in proposals are similarly punished by moving to s^{D_1} or s^{D_2}.

To verify that the norm is self-sustaining, notice that whenever 1 gets to make a proposal (at any of continuations considered in the construction), his payoff is at least $\pi_1(D_1)$. Since D_1 is on the Pareto frontier, any proposal giving 1 strictly more than $\pi_1(D_1)$ must give 2 less than $\pi_2(D_1)$. This means that s^{D_1} can serve as a punishment in case 1 makes such a proposal.

Remaining cases: $\underline{v}_1 < \pi_1(P_2)$ and/or $\underline{v}_2 < \pi_2(P_1)$

The construction is almost identical in other cases. The only difficulty is that the difference $\pi_1(A_1) - \underline{v}_1$ is now bounded below away from zero, whereas it was previously of order $\sqrt{\varepsilon}$. This may lead to situations in which the points C_{1j} constructed above are no longer feasible and/or give 2 a payoff lower than $\pi_2(A_2)$. The difficulty is easily addressed by adding, for each j, a point E_{1j} lying on the segment $[A_1B_{1j}]$—and thus also on the line (A_1C_{1j})—such that if player 2 gets a chance to make a proposal, or if nobody does, players' continuation payoffs jump to E_{1j}. The promise keeping condition (6) becomes

$$\pi_1(B_{1j}) = \left(1 - \frac{p}{2}\right)\pi_1(E_{1j}) + \frac{p}{2}\pi_1(C_{1j}) \tag{9}$$

[38] It is possible to show that each of A, A_1, and A_2 can be implemented so that players' continuation payoffs eventually converge to a Pareto-efficient point. Under this "redemptive" implementation, if players switch to a Pareto-inefficient element following a deviation, they will eventually forgive and forget past deviations.

Choosing E_{1j} close enough to B_{1j} ensures that C_{1j} lies within a distance $\sqrt{\varepsilon}$ of B_{1j} and, hence, of A_1. This guarantees that C_{1j} is feasible and does not drop below $\pi_2(A_2)$, so that the rest of the argument for the first case can be applied. To implement $s^{E_{1j}}$, we use public randomization to implement it as a probabilistic mixture of s^{A_1} and $s^{C_{1j}}$.

A.2 Necessary Conditions

The interesting case is when $\underline{v}_i < \pi_i(P_{-i})$: otherwise, Theorem 1 predicts only that i's payoff must be individually rational. We derive the necessary condition for player 1; the same argument can be applied to player 2.

Suppose that $\pi_1(P_2) > \underline{v}_1$ and, by contradiction, that there is a point $A \in V(q)$ such that $\pi_1(A) < v_1 = \underline{v}_1 + \frac{q}{2+q}(\pi_1(P_2) - \underline{v}_1)$: one can construct, for any ε small enough and per-period probability $p = q\varepsilon$ of proposal opportunity, a self-sustaining norm s that implements A.

Let C_1 denote 1's infimum payoff over all continuations of s following histories at which it is 1's turn to make a proposal and no off-path proposal has yet been accepted. Since the Pareto point P_2 is a possible proposal payoff,[39] and since it Pareto dominates all payoffs with $\pi_1 < \pi_1(P_2)$, C_1 must satisfy $\pi_1(P_2) \leq C_1$.

We now contradict this inequality. Let $\mathcal{N}$ denote the set of continuations of s at the beginning of all periods following histories at which no off-path proposal has been accepted. Also let $A_1 = \inf_{V \in \mathcal{V}(\mathcal{N})} \pi_1(V)$, $B_1 = \inf_{W \in \mathcal{W}(\mathcal{N})} \pi_1(W)$, and $D_1 = \inf_{U \in \mathcal{U}(\mathcal{N})} \pi_1(U)$, and consider any sequence $\{V_k\} \in \mathcal{V}(\mathcal{N})$ such that $\pi_1(V_k) \to_{k \to +\infty} A_1$. For any V_k there is an action that implements it in the first period of the corresponding continuation. However, if player 1 deviates, he can guarantee himself an immediate payoff of at least $\underline{v}_1$, and the worst punishment for him after deviation gives him at least B_1. Therefore, $\pi_1(V_k) \geq \varepsilon \underline{v}_1 + (1-\varepsilon)B_1$. Since this inequality holds for all V_k we obtain, taking the limit:

$$A_1 \geq \varepsilon \underline{v}_1 + (1-\varepsilon)B_1 \tag{10}$$

Since any element of $\mathcal{U}(\mathcal{N})$ lies in the convex hull of $\mathcal{V}(\mathcal{N})$, and player 1 can always conceal his opportunity to propose, we have $C_1 \geq D_1 \geq A_1$. Consider now a sequence $\{W_k\} \in \mathcal{W}(\mathcal{N})$ such that $\pi_1(W_k) \to B_1$. Any element W_k is a weighted average of an expected payoff vector EU_k^1 whenever 1 gets a chance to make a proposal, an expected payoff vector EU_k^2 when it is 2's turn to make a

[39] By the Folk Theorem, P_2 can be implemented by an equilibrium of the repeated game without renegotiation. P_2 can thus also be implemented as a norm of the enlarged game in which all proposals are ignored, i.e., treated as cheap talk.

proposal, and a payoff vector U_k^0 in case no one gets to make a proposal:

$$W_k = \frac{p}{2}(EU_k^1) + \frac{p}{2}(EU_k^2) + (1-p)(U_k^0) \tag{11}$$

We note that EU_k^1 is a mixture of elements of $\mathcal{U}(\mathcal{N})$ resulting from 1's mixture over proposals and 2's mixture over her acceptance decision. Similarly, EU_k^2 is a mixture of elements of $\mathcal{U}(\mathcal{N})$.

Since all elements U_k's belong to $U(\mathcal{N})$, we have $\pi_1(EU_k^2) \geq A_1$ and $\pi_1(U_k^0) \geq A_1$. Equation (11) thus implies that

$$\pi_1(W_k) \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}\pi_1(EU_k^1).$$

Recalling that C_1 denotes 1's infimum payoff when he gets to make a proposal, we get

$$\pi_1(W_k) \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}C_1.$$

Taking limits, $B_1 \geq (1 - \frac{p}{2})A_1 + \frac{p}{2}C_1$, or

$$B_1 \geq (1 - \frac{q\varepsilon}{2})A_1 + \frac{q\varepsilon}{2}C_1. \tag{12}$$

Combining (10) and (12), we conclude that $A_1 \geq \varepsilon\underline{v}_1 + (1-\varepsilon)[(1 - \frac{q\varepsilon}{2})A_1 + \frac{q\varepsilon}{2}C_1]$. Ignoring terms of order ε^2 in the right-hand side of this equation, $A_1 \geq \varepsilon\underline{v}_1 + (1 - [1 + \frac{q}{2}]\varepsilon)A_1 + \frac{q\varepsilon}{2}C_1$. Subtracting A_1 on both sides of the last equation and dividing by ε, we obtain

$$0 \geq \underline{v}_1 - [1 + \frac{q}{2}]A_1 + \frac{q}{2}C_1 \tag{13}$$

From $A_1 \leq \pi_1(A)$, $C_1 \geq \pi_1(P_2)$, and $\pi_1(A) < v_1 = \underline{v}_1 + \frac{q}{2+q}(\pi_1(P_2) - \underline{v}_1)$, we get

$$0 < \underline{v}_1 - [1 + \frac{q}{2}]A_1 + \frac{q}{2}C_1$$

which contradicts (13). This shows the necessary condition for player 1.

An identical reasoning for player 2 shows the second necessary condition. This proves the result for $P_1 \neq P_2$. A similar reasoning applies when $P_1 = P_2$.

Credible proposals Section 5 has introduced the concept of $\mathcal{C}$-credible proposals, and claimed that the necessity conditions were unaffected if the proposals involved in the definition of stability were restricted to being credible. To prove this claim, it suffices to verify that the proposal to move to P_2, used just above to derive the necessary condition, is $\mathcal{C}$-credible. The norm implementing P_2 is constructed as follows: players are prescribed to play, in all periods, the pure action profile with payoff P_2, and to abstain from making any proposal. Any deviation, whether in action or in proposal, triggers the continuation implementing A—which is supposed to exist, by the contradiction hypothesis. Clearly, player 2 cannot benefit from deviating as she is getting her highest possible payoff in the game. Moreover, the difference $\pi_1(P_2) - \pi_1(A)$ is by assumption bounded below by $\frac{2}{2+q}(\pi_1(P_2) - \underline{v}_1)$, which is ε-independent. Therefore, 1 cannot benefit from deviating either: a deviation in action may create an immediate gain of order ε, but triggers a drop in continuation payoffs that is ε-independent and dominates the gain. A deviation in proposal yields the payoff vector A, which again is detrimental to 1.

B Concept equivalence: Proof of Theorem 2.

1. Any closed convention $\mathcal{C}^c$ is an open convention as well, so the first statement is trivially true. Now consider any open convention $\mathcal{C}^o$. To construct a payoff-equivalent closed convention $\mathcal{C}^c$, we modify each norm s of $\mathcal{C}^o$ as follows: s's rules on and off the equilibrium path are kept unchanged except when a player, say i, sends a message m_i which is off the equilibrium path. In this case, because $\mathcal{C}^o$ is an open convention, the continuation if $-i$ accepts the proposal need not lie in $\mathcal{C}^o$. Following such a proposal, players are instead prescribed to behave as if i had remained silent. The new rules define a norm: when playing the original norm s, i was not sending the message m_i anyway, so removing this option does not affect equilibrium behavior and payoffs. By construction, the set of modified norms form a closed convention $\mathcal{C}^c$, and because each norm of $\mathcal{C}^o$ has been modified into a single payoff-equivalent norm of $\mathcal{C}^c$, the conventions are payoff equivalent.

2. We start with the observation that if two conventions $\mathcal{C}^c$ and $\mathcal{C}^o$ have the same payoff sets, then any proposal that is credible according to either convention is credible according to the other one.

We now consider any stable open convention $\mathcal{C}^o$ and construct the corresponding closed convention $\mathcal{C}^c$ as in Part 1. To show that $\mathcal{C}^c$ is stable, consider any norm s of $\mathcal{C}^c$, history at which player i gets to propose, and credible proposal U such that $\pi_i(U)$ is strictly greater than i's continuation payoff $\hat{U}_i$. From the above observation, U is also credible for $\mathcal{C}^o$. If the proposal U gives player $-i$ a lower payoff that $\hat{U}$ does, then the payoff $U' = \hat{U}$ satisfies Definition 5. If the proposal U Pareto dominates $\hat{U}$, then for the norm $\tilde{s}$ of $\mathcal{C}^o$ corresponding to s, and the same history, $-i$ must reject U with positive probability (for otherwise $\pi_i(U)$ would coincide with $\hat{U}_i$). Let U' denote the continuation payoff if $-i$ rejects U. By stability of $\mathcal{C}^o$, $-i$ knows that if he accepts U it will be implemented. Since it is weakly optimal for $-i$ to reject U, it must therefore be the case that $\pi_{-i}(U') \geq \pi_{-i}(U)$. Moreover, it must also be the case that $\pi_i(U') \leq \hat{U}_i$, for otherwise it would be strictly optimal for i to deviate by proposing U, and $\tilde{s}$ would not be a part of an open stable convention $\mathcal{C}^o$. Using this U' in Definition 5, this implies that $\mathcal{C}^c$ is stable.

Next, consider any stable closed convention $\mathcal{C}^c$. To construct a payoff-equivalent stable open convention $\mathcal{C}^o$, we simultaneously modify all norms of $\mathcal{C}^c$. The modification proceeds in two steps, using a recursive representation norms. A norm may be viewed as a prescription of actions, proposals and acceptance decisions for the next period (each depending on what happened in earlier stages), along with a continuation norm resulting from these stages applied to the period after next. In Step 1, we modify the prescriptions for time $t+1$, and still use norms of $\mathcal{C}^c$ as continuation norms. The purpose of this step is to make a prescription compatible with the requirement that if a Pareto-improving, credible proposal is made and accepted, then it has to be played. In Step 2, we replace these continuation norms of $\mathcal{C}^c$ by those built in Step 1, to guarantee that the rule applies at all periods, ensuring that credible norms which are accepted are implemented, so that Definition 7 holds at all periods.

Consider any norm s of $\mathcal{C}^c$. We modify s as follows. For the modified norm $\tilde{s}$, the action stage and on-path proposals are prescribed exactly as in s.[40] Now consider a history at which i makes any proposal U which is not prescribed by s but which is $\mathcal{C}^c$-credible. If $-i$ accepts the proposal, we construct $\tilde{s}$ by prescribing that players implement this proposal.[41] If the proposal gives i a strictly higher payoff than his continuation payoff $\hat{U}_i$, then by stability of $\mathcal{C}^c$, there must exist a payoff vector U' corresponding to some norm s' of $\mathcal{C}^c$, which gives player $-i$ at least as much as U, and which gives player i at most $\hat{U}_i$. We prescribe playing the norm corresponding to U' in case player $-i$ rejects the proposal. If U does not improve upon i's continuation payoff, we prescribe playing the continuation corresponding to any of i's on-path proposals in case $-i$ rejects U. Finally, if i makes a non-credible proposal, the proposal is ignored as if i had stayed silent.

We now verify that $\tilde{s}$ is a norm that yields the same payoff as s. Since $\tilde{s}$ prescribes the same actions as s, players are incentivized to follow the prescription. If i gets a chance to make a proposal, any proposal prescribed by s (and hence $\tilde{s}$) yields the same continuation payoff as in s. If player i makes a credible, off-equilibrium proposal that improves upon his on-path payoff, then player $-i$ is incentivized to reject it, and i's continuation payoff is weakly lower than his on-path payoff. It is never optimal for i to make a credible proposal that is lower than his on-path payoff, regardless of $-i$'s acceptance decision. Finally, we replace all continuation norms by their modified versions.

[40] Another modification of s is needed when i proposes on path a continuation $\hat{s}$ that lies outside of $\mathcal{C}^c$, which $-i$ is supposed to accept, and which is followed by a continuation s' in the convention $\mathcal{C}^c$ (as it should, since the convention is closed). This sequence of moves is replaced by i directly proposing s' and having it accepted by $-i$. The modified profile is also a norm, as is easily checked. More generally, any norm of the game may be turned into a payoff-equivalent "truthful" norm of the game, i.e., one in which any proposal that is made and accepted *on the equilibrium path* is implemented, as explained and proved in Appendix F.

[41] At this point, we do not know yet that the proposal is $\mathcal{C}^o$-credible. We only know that it is $\mathcal{C}^c$-credible. However, the norm $\mathcal{C}^o$ that we are constructing will be payoff equivalent to $\mathcal{C}^c$ and hence have the same set of credible proposals.

There remains to verify that the set consisting of all modified norms forms a stable open convention, denoted $\mathcal{C}^o$, which is payoff equivalent to $\mathcal{C}^c$. First, we notice that continuations outside of $\mathcal{C}^o$ may arise only when a player makes an off-path proposal (which, by construction, also has to be credible) which is accepted by the other player. Thus, $\mathcal{C}^o$ is an open convention. By construction, each element of $\mathcal{C}^o$ corresponds to exactly one element of $\mathcal{C}^c$, which yields the same expected payoff. Therefore, the conventions are payoff equivalent. As observed earlier, this implies that they have the same set of credible proposals. This, in turn, implies that any Pareto-improving, credible proposal of $\mathcal{C}^o$ that is accepted is played and, hence, that $\mathcal{C}^o$ is stable.

C Proofs of Section 6 (Arbitrary number of players)

C.1 Proof of Proposition 3

We fix one of the two procedures and let $\mathcal{F}_k$ denote the set corresponding to the k-th step in the sequential reduction of the set $\mathcal{F}$ under this procedure. We first show that points on the relative Pareto frontier $\mathcal{P}(\mathcal{F}_k)$ of $\mathcal{F}_k$ are never removed by the procedure. Suppose, contrary to the claim, that some point $A \in \mathcal{P}(\mathcal{F}_k)$ was removed by the procedure. Then there would be a player i such that $\pi_i(A) < \underline{\pi_i}(\mathcal{F}_k)$. If A was prescribed as a punishment payoff for any proposal of player i, then for i's optimal proposal with payoff $C \in \mathcal{F}_k$, the punishment payoff A would not be credible as it is removed at the k-th step. That is, any $j \neq i$ has $\pi_j(A) < \pi_j(C)$. Since A lies on the Pareto frontier of $\mathcal{F}_k$, this means that $\pi_i(C) < \pi_i(A)$: C gives i a lower payoff than $\underline{\pi_i}(\mathcal{F}_k)$, which contradicts C's assumed optimality. One could simply prescribe both continuations to have C as their payoff vector, and this would give i a lower payoff than $\underline{\pi_i}(\mathcal{F}_k)$.

When evaluating the worst punishment $\mathcal{D}(C, \mathcal{F}_k)$ for player i for making a (non-prescribed) proposal with payoff C, the optimal proposal (that is, the one which gives the highest payoff to player i from the worst punishment) always lies on the Pareto frontier. Indeed, consider a proposal with payoff C, which is not Pareto-optimal, and another proposal with payoff C', which Pareto dominates C. The set $\mathcal{D}(C, \mathcal{F}_k)$ of possible punishment payoffs is strictly larger than the set $\mathcal{D}(C', \mathcal{F}_k)$, since the latter set gives every player $j \neq i$ a higher lower-bound on his payoff. This implies that the proposal C' gives player i a worst punishment payoff $\underline{\pi_i}(C', \mathcal{F}_k)$ at least as high as the proposal associated with payoff C.

Since no point on the relative Pareto frontier of $\mathcal{F}$ is removed in the sequential reduction, the set of optimal proposals (in terms of evaluating the worst possible punishment) for any player i remains the same along the sequence. However, the set of possible punishments keeps decreasing at each step, which weakly increases, as a result, the minimal value $\underline{\pi_i}(\mathcal{F}_k)$ with k. (Recall that $\underline{\pi_i}(\mathcal{F}_k)$ is i's minimal payoff if he gets a chance to make a proposal). At each step, the set $\mathcal{F}_k$ is characterized by the n lower bounds of the players' payoffs $\{\pi_{min,i}(\mathcal{F}_k)\}_{i\in\{1,\ldots,n\}}$. These lower bounds are weakly increasing at each step, which implies that the procedure converges to a stable point.

C.2 Proof of Theorem 4

Necessity

Suppose that A lies outside of $\mathcal{V}_N$ and, for any $\varepsilon > 0$ small enough, there exists a self-sustaining norm $s(\varepsilon)$ such that A lies in the set $\mathcal{V}(\varepsilon)$ of payoff vectors across all continuations of $s(\varepsilon)$ following histories at which no off-path proposal was accepted by the supermajority. The sets $\mathcal{V}(\varepsilon)$ must satisfy inequality (3) (replacing $\mathcal{F}_N^k$ as an argument of this inequality), up to an ε-term. Consider the limit of $\mathcal{V}(\varepsilon)$ as ε goes to 0. This limit payoff set contains A and satisfies inequality (3), which implies that A should have not been removed from any of the sets $\mathcal{F}_N^k$. However, this implies that A belongs to $\mathcal{V}_N$, a contradiction.

Sufficiency: $\mathcal{R}$

We first prove that any point in $\mathcal{R}$ is sustainable. Consider any point A with $\pi_i > \pi_i(P_{-i})$ for any i. As in the two-player case, one can find n points A_i such that for $j \neq i$ $\pi_j(A_i) = \pi_j(A)$ and $\pi_i(A_i) = \pi_i(P_{-i}) + \sqrt{\varepsilon}$. We build a self-sustaining norm s, in which for each i, the point A_i gives i his lowest payoff across all continuations, following histories at which no off-path proposal was accepted. In the continuation norm s^{A_i} associated with payoff vector A_i, player i is being minmaxed. Since players other than i may have to use mixed strategies, this generates a set $\mathcal{B}$ of continuation payoffs, following the action stage, which depend on the realization of actions of players other than i. Any continuation $B \in \mathcal{B}$ is implemented as follows: if i can make a proposal, he is prescribed to propose some continuation with payoff C; other players are prescribed to remain silent; in the absence of any proposal, the continuation returns to s^{A_i}. As in the two-player case, one can guarantee (possibly using the public randomization), that the distance A_iC is of order $\sqrt{\varepsilon}$.

Since the Pareto frontier is connected, so is its truncation to points for which i's payoff lies above $\pi_i(P_{-i})$. One can therefore find a connected subset S_ε of the frontier consisting of all points giving, for each i, a payoff greater than or equal to $\pi_i(A_i) + K\varepsilon$, where K is a constant chosen large enough that players are incentivized not to deviate in actions.

Continuation norms with payoffs in S_ε are constructed in such a way that each player i gets at least $\pi_i(A_i) + K\varepsilon$ in all subsequent continuations.

When implementing A_i, players are already incentivized to follow the prescribed actions. If i makes a non-prescribed proposal, then by construction of S_ε there exists a continuation with a payoff Q_i in S_ε that gives i a lower payoff than C. Indeed, the lower bound for π_i at the set S_ε is $\pi_i(A_i)+K\varepsilon$, while $\pi_i(C) - \pi_i(A_i)$ is of order $\sqrt{\varepsilon}$.

Sufficiency: General Conditions

The proof is similar to that of the two-player case. For any point $A \in \mathcal{V}_S$ with $\pi_i > \pi_{min,i}(\mathcal{V}_S)$, consider the set of points $A_i \in \mathcal{V}_S$ such that for any i $\pi_i(A_i) = \pi_{min,i}(\mathcal{V}_S) + \sqrt{\varepsilon}$ and $\pi_{-i}(A_i) = \pi_{-i}(A)$. The points A_i have a smaller i-th coordinate than A provided that ε is small enough. We also assume without loss of generality that $\pi_i(A_j) - \pi_i(A_i) >> \sqrt{\varepsilon}$ for any $j \neq i$.

We build a self-sustaining norm s such that A_i gives the lowest payoff to player i across all continuations following histories at which no off-path proposal was accepted. At s^{A_i}, player i is minmaxed. Since players other than i may have to mix their actions, we construct a set of continuations with payoffs $B \in \mathcal{B}$, corresponding to the observed actions of players $-i$. For any continuation norm s^B associated with some payoff $B \in \mathcal{B}$, i is prescribed to make a proposal with some payoff vector C, and all other players are prescribed to remain silent. As with the two-player case, C can be assumed to lie at a distance of order $\sqrt{\varepsilon}$ from A_i. When implementing the continuation norm s^C associated with C, players are prescribed to follow a deterministic sequence of actions such that the continuation payoff remains within an ε-distance from C. Players are prescribed not to make any proposals.

The initial point A is also implemented by deterministic actions and no proposals. Moreover, each point in the positive orthant starting at the vertex with i^{th} coordinate $\pi_i(A_i) + K\varepsilon$ for each i is implemented by a self-sustaining continuation norm of s in such a way that $\pi_i > \pi_i(A_i) + K\varepsilon$: s^{A_i} is a severe enough punishment for i that it is suboptimal for him to deviate in action.

The norm s can be shown to be generically self-sustaining. The only new issue concerns i's incentives to deviate in proposal. We have reduced (increased the lower bounds on payoffs) the initial set $\mathcal{V}_S$ by an order of $\sqrt{\varepsilon}$. The orthant defined by $\pi_i > \pi_i(A_i) + K\varepsilon$ for all i is part of the set of self-sustaining continuations, but some points lying below this orthant are removed from the original set $\mathcal{V}_S$. As a result, the value $\underline{\pi_i}(.)$, which i can guarantee if having a chance to propose, can now be larger. Our goal is to show that, nevertheless, generically the value of $\underline{\pi_i}(.)$ is smaller than $\pi_i(C)$, and therefore player i is incentivized to propose s^C.

When building a set $\mathcal{V}_S$ by sequentially removing payoffs with $\pi_{min,i}(.) < \underline{\pi_i}(.)$, the initial set of individually-rational payoffs gets reduced. If for player i the final value of $\pi_{min,i}(\mathcal{V}_S)$ is strictly larger than his minmax payoff $\underline{v_i}$, then the value of $\pi_i(A_i) - \underline{v_i}$ is of order ε^0. This means that the distance A_iC can be made of $\varepsilon^{\frac{1}{4}}$-order. At the same time, the set $\mathcal{V}_S$ (and, respectively, the value $\underline{\pi_i}(.)$) were changed by an order of $\sqrt{\varepsilon}$, guaranteeing that $\underline{\pi_i}(.) < \pi_i(C)$.

If player i's payoff $\pi_{min,i}(\mathcal{V}_S)$ equals to minmax $\underline{v_i}$, this means that i's payoff was not increased when building the set $\mathcal{V}_S$. Put it differently, one can consider a hyperplane of the set $\mathcal{V}_S$ with $\pi_i = \underline{v_i}$, and find the maximum payoffs of other players $\overline{\pi_j}$, $j \neq i$ on that hyperplane. The $n-1$-dimensional payoff vector $\{\overline{\pi_j}\}_{j\neq i}$ cannot lie within the interior of $\mathcal{V}_S$ (otherwise, player i could make a proposal dominating $\{\overline{\pi_j}\}_{j\neq i}$ and thus guaranteeing himself a payoff higher than $\underline{v_i}$). When the set $\mathcal{V}_S$ is reduced by (an arbitrarily small) $\sqrt{\varepsilon}$-order, player i can gain incentives to make an off-path proposal, only if the vector $\{\overline{\pi_j}\}_{j\neq i}$ lies exactly on the Pareto frontier of $\mathcal{V}_S$. However, this possibility is non-generic.

C.3 Proof of Proposition 4

Intuition. The sets $\mathcal{V}_S$ and $\mathcal{V}_N^q$—necessary conditions depend on q, hence the superscript—are both obtained from $\mathcal{F}$ by sequentially increasing the lower bounds on each player's payoff when he gets a chance to make a proposal. $\mathcal{V}_S$ is obtained by removing payoffs below $\underline{\pi_i}(\cdot)$ at each step, while $\mathcal{V}_N^q$ is obtained by removing payoffs below $\frac{n\underline{v_i}+q\underline{\pi_i}(.)}{n+q}$. When q goes to infinity, the sets of payoffs removed at each step of these procedures converge to each other. As we show below, this implies that $\mathcal{V}_N^q$ converges to the set $\mathcal{V}_S$ as q goes to infinity.

The set of sufficient conditions, $\mathcal{V}_S$, can be characterized by two sets of lower bounds for each player i: $\underline{\pi_i}(\mathcal{V}_S)$ is the lower bound on i's payoff when he gets a chance to make a proposal and $\pi_{min,i}(\mathcal{V}_S)$ is the lower bound for his payoff at the beginning of a period. $\mathcal{V}_S$ was constructed in such a way that $\underline{\pi_i}(\mathcal{V}_S) \leq \pi_{min,i}(\mathcal{V}_S)$.

To capture this intuition, we first show by induction that $\mathcal{V}_S$ is the largest set $\mathcal{S}$ of individually rational payoffs whose Pareto frontier is equal to $\mathcal{P}(\mathcal{V})$ and such that $\underline{\pi_i}(\mathcal{S}) \leq \pi_{min,i}(\mathcal{S})$ for any i. Consider such a set $\mathcal{S}$. The sequence of sets $\mathcal{F}_S^k$ converging to $\mathcal{V}_S$ starts with $\mathcal{F}_S^0 = \mathcal{F}$, the set of all individually rational points. This implies that $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_S^0)$, since $\mathcal{F}_S^0$ contains $\mathcal{S}$ and, hence, the set of punishments if i makes an unprescribed proposal is higher with $\mathcal{F}_S^0$ than with $\mathcal{S}$, resulting in a lower bound $\underline{\pi_i}$. We now show the induction hypothesis: if $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_S^k)$, then the same condition holds for $k+1$. Due to the way the payoffs are cut at step k, one has for each i, $\pi_{min,i}(\mathcal{F}_S^{k+1}) = \max\{\pi_{min,i}(\mathcal{F}_S^k), \underline{\pi_i}(\mathcal{F}_S^k)\} \leq \underline{\pi_i}(\mathcal{F}_S^k)$, which does not exceed $\underline{\pi_i}(\mathcal{S}) \leq \pi_{min,i}(\mathcal{S})$. Since the lower bound $\pi_{min,i}(\mathcal{F}_S^{k+1})$ is lower than $\pi_{min,i}(\mathcal{S})$, the set $\mathcal{F}_S^{k+1}$ contains $\mathcal{S}$, and one has that $\underline{\pi_i}(\mathcal{S}) \geq \underline{\pi_i}(\mathcal{F}_S^{k+1})$. By induction, the limit set $\mathcal{V}_S$ contains $\mathcal{R}$.

Let $\mathcal{V}_N$ denote the limit of $\mathcal{V}_N^q$ as q goes to infinity. We wish to show that $\mathcal{V}_N = \mathcal{V}_S$. Consider the sequences $\{\mathcal{F}_N^{k,q}\}_{k=0}^{+\infty}$ resulting from the procedure applied, for any fixed q, to derive necessary conditions for this value of q. Due to the way points are removed at each step, it is easy to check that $\mathcal{F}_N^{k,q'} \subset \mathcal{F}_N^{k,q}$ whenever $q' > q$; by the same logic, it is straightforward to check that $\mathcal{V}_S$ is contained in $\mathcal{V}_N$. To prove the reverse inclusion, note for each q and i, we have $\pi_{min,i}(\mathcal{V}_N^q) \geq \frac{n\underline{v_i}+q\underline{\pi_i}(\mathcal{V}_N^q)}{n+q}$, as this inequality holds at each step k of the procedure. Taking the limit as q goes to infinity, the limiting set $\mathcal{V}_N$ must satisfy for each i $\underline{\pi_i}(\mathcal{V}_N) \leq \pi_{min,i}(\mathcal{V}_N)$. From the previous paragraph, this implies that $\mathcal{V}_S$ contains $\mathcal{V}_N$, which concludes the proof.

For Online Publication

From Self-Enforcing Agreements to Self-Sustaining Norms

Mikhail Safronov
University of Cambridge

Bruno Strulovici
Northwestern University

December 23, 2017

*We are grateful for comments from David Miller, Larry Samuelson, Ilya Segal, Tadashi Sekiguchi, Takuo Sugaya, Joel Watson, and seminar participants at the 2015 SAET meeting, the 2015 World Congress of the Econometric Society, and the 2016 Asian Meeting of the Econometric Society. Strulovici gratefully acknowledges financial support from an NSF CAREER Award (Grant No. 1151410) and a fellowship from the Alfred P. Sloan Foundation

D Messages as equilibria: A conceptual difficulty

In the main text, we defined the negotiation stage through an exogenously given set of messages, and then assigned a norm to each message, requiring that each payoff vector of the feasible set be assigned a message. One may wonder if it is possible to assign a message to each *norm*, rather than to each payoff. This requirement is much stronger since there are potentially many equilibria that yield the same payoff. Moreover, this creates a circularity problem, because if one enlarges the message space, this changes the underlying game (since more messages can be sent) and thus the set of equilibria.

To be specific, suppose that the description of each period, in Section 2, is modified as follows. Steps 1–3 and 5 are unchanged, but Step 4 is replaced by the following step:

4') With probability $p < 1$, one of the players is chosen to propose a new *plan* for continuation of the game. Each player has the same probability of $\frac{p}{2}$ being chosen. The chosen player may conceal his proposal opportunity by remaining silent, or mix between making a proposal or staying silent.

The object of a proposal is an infinite-horizon plan m from the set $\mathcal{M}$ of all possible plans, defined as follows

A *plan* in period t describes players' strategy for the infinite repetition of the stage-game described above, *from period* $t+1$ *onwards*. These decisions (actions, proposals, and acceptance mixtures) are history-dependent. The setting being time invariant, it is convenient to define recursively the set $\mathcal{M}$ of plans. A plan $m \in \mathcal{M}$ in period t is described by the following elements:

a) For each realization z of the public randomization device, a pair $\alpha = \alpha[m](z)$ of mixed actions that players should play in period $t+1$;

b1) For each player i, a distribution $\bar{\mu}_i = \bar{\mu}_i[m](z, \mathbf{a}) \in \Delta(\mathcal{M} \cup \emptyset)$ over proposals, where the outcome $\emptyset$ means that i abstains from making a proposal (unbeknownst to player $-i$). We assume that distributions have a finite support over plans.[1] The proposal distribution can depend on the realization z of the public randomization device and on the pair $\mathbf{a}$ of observed actions. Because $p < 1$, not observing any proposal from either player is always consistent with "on-path" behavior. The realized proposal is denoted μ_i;

b2) A probability $q_{-i} = q_{-i}[m](z, \mathbf{a}, \mu_i)$ that $-i$ accepts i's proposal (whenever $\mu_i \neq \emptyset$), which may depend on z, $\mathbf{a}$, and μ_i;

[1] We will in fact impose a uniform upper bound on this support, as explained below.

b3) If no one has made a proposal, the acceptance stage is skipped. To economize on notation, we assume that some player i is, even in that case, conventionally selected (randomly or deterministically) as the proposer and let $\mu_i = \emptyset$ and $D_{-i} = 0$. (So, $-i$'s conventional response is to systematically "reject" a non proposal.)

c) A continuation plan $m_{+1} = m_{+1}[m](z, \mathbf{a}, i, \mu_i, D_{-i}) \in \mathcal{M}$ for period $t+2$ onwards, as a function of z, $\mathbf{a}$, i, μ_i, D_{-i}, where i indicates the identity of the last proposer.[2]

While the above definition seems natural, it turns out to be too permissive for the set of plans to be well-defined: there does not exist a set of plans so large as to contain all the possible continuation prescriptions allowed above. In particular, in the above construction, a plan must specify an acceptance decision for each possible proposal. Therefore, each plan m must specify—among other things—a function which maps each element of $\mathcal{M}$ (the proposal) to a binary decision (acceptance). This implies that the set $\mathcal{M}$ of plans must contain, in order to include all possible prescriptions, its power set $2^{\mathcal{M}}$. Such a set does not exist, since any set has a strictly lower cardinality than its power set, by Cantor's Power Set Theorem (see, e.g., Mendelson (1997)).

E Asymmetric proposing probabilities

It is easy to extend the analysis to a protocol in which one of the players has a higher probability factor q_i of proposal than the other player. The sufficient conditions are unchanged in this setting, but the necessary conditions become tighter for the player whose proposal probability is higher, which translates into a higher minimal guaranteed payoff for that player across all self-sustaining norms. To see this clearly suppose that $\underset{\sim}{v}_1 < \pi_1(P_2)$ and $\underset{\sim}{v}_2 < \pi_2(P_1)$ (configuration (a) in Figure 2), so that renegotiation potentially benefits both players, compared to the minmax payoffs, and consider the case in which 1 can make frequent proposals while 2 never gets a chance to make a proposal (i.e., q_1 is arbitrarily large while $q_2 = 0$). Then, 2's minimal guaranteed sustainable payoff collapses to her minmax payoff, while 1 is guaranteed to get a payoff of at least $\pi_1(P_2)$. More generally, player i's minimal payoff, given by (2), is calculated using the probability q_i that he gets an opportunity to make a proposal, and is independent of the other player's probability of getting that opportunity. As q_i increases, player i's guaranteed continuation payoff increases as well, and vice versa.

[2]Clearly, this plan must be independent of i whenever $\mu_i = \emptyset$, so that the convention chosen for the proposer in the absence of any actual proposal is indeed irrelevant. This restriction is applied throughout.

F Truthful norms

For any number of players and any norm there is a payoff-equivalent norm, which is *truthful* in the sense that any on-path proposal is always accepted and implemented. Indeed, when some player i gets a chance to make a proposal, he can make any number of proposals in equilibrium, the expectation of which is some continuation payoff C. We alter the norm by prescribing player i to make only one proposal with payoff C. The altered norm prescribes all other players to accept the proposal and C to be implemented regardless of the acceptance decision. The payoff C can be implemented using public randomization.

If i deviates and proposes a Pareto improvement relative to C, everyone is prescribed to reject it.

With two players, the new norm prescribes to have the same rejection continuation as in the original norm. The incentive to accept the proposal is unaffected by the transformation, so the other player is incentivized to reject an off-path proposal.

With more than two players rejecting the off-path proposal is an equilibrium. When the norm is simple (see Section 6) and players vote for the payoff-improving proposal, as in Assumption 1, player i is still prescribed to propose C. If player i makes an off-path proposal which gives him more than $\pi_i(C)$, for each such a proposal there is at least one player $j \neq i$ who rejects it, as otherwise this off-path proposal would be made and accepted in the original norm.

G General renegotiation protocols

The benchmark model can be extended to allow multiple rounds of renegotiation within each period without affecting the main results. We consider a multi-round renegotiation environment similar to Miller and Watson (2013), adopting their notation. In each period, with probability $q_i\varepsilon$ player i can make a proposal to player j, which initiates a stochastic alternating-offer renegotiation protocol. The renegotiation rounds are numbered as $l \in \{1, 2, ...\}$,with $l = 1$ being the original proposal of player i. If at any round the proposal is accepted, the players stop renegotiation and move to the next period; otherwise the players continue renegotiation. Conditional on reaching the round $l > 1$, player i is selected to be the proposer with probability $\rho_{i,l}$. Conditional on rejecting the proposal in l-th round, the renegotiation breaks with probability $\beta_l \in [0, 1]$, with $\prod_{l=1}^{\infty}(1 - \beta_l) = 0$. The values of $\rho_{i,l}$, $l > 1$, and of β_l, $l \geq 1$, are assumed to be independent of made proposals, and of the identity of the original proposer, and are the same across different time periods.

The presence of multiple rounds of renegotiation affects the set of equilibria in the repeated game with renegotiation. Nevertheless, Definitions 2, 3, **??** have straightforward extensions applied to each round of renegotiation and Theorem 1 continues to hold as stated. It is equally easy to show for environment of Section L, below, where the continuation of a failed proposal is independent of the proposer, that Theorems 8 and 9 also continue to hold.[3]

THEOREM 5 *Theorem 1 extends to multi-round renegotiation.*

The proof for sufficient conditions in Theorem 1 still holds, since player i can be punished for making an unprescribed proposal by moving to the continuation which gives the highest possible payoff to the player j, which makes all the future rounds of renegotiation meaningless. The proof for the necessary conditions in Theorem 1 also holds, since when player i gets a chance to propose, he can guarantee to move to the above continuation, in both the one-round and the multi-round cases.

H Comparative statics

Consider any $q > q'$ and any norm s that is self-sustaining at some frequency q. We will show the existence of a norm s' that is self-sustaining at frequency for q' and payoff-equivalent to s.

Under the new norm s', any payoff A achieved by s following any history ending before the action stage is implemented using the same mixed actions and the same subsequent continuations as prescribed by s. Consider now any vector payoff B, calculated before the proposal stage, implemented by some continuation s^B of s. s^B is a mixture of three continuation equilibria: s^{C_1}, which arises when 1 gets a chance to make a proposal and is calculated after the proposal stage; s^{C_2} which arises if 2 gets to make a proposal; and s^C, which arises if no one gets to make a proposal.

At frequency q', B is implemented as follows: players are prescribed to make exactly the same proposals (with the same prescribed punishments if someone made an off-path proposal). For B to still to be the weighted average of the continuations occurring after the three proposal events, we change the continuation payoff in case no proposal is made: the new continuation payoff in this case, C', has to lie on the line between B and C. The new continuation $s^{C'}$ is achieved using public

[3]Concerning Theorem 8, even in the multi-round case renegotiation ends in one round, with a continuation payoff being on the Pareto frontier. Moreover, with each player i having a non-zero bargaining power (equivalent to $\pi_i > 0$), the proof of Theorem 9 in Section L.1 does not require the property of η-stability, since the Pareto-improving proposal will lie in the interior of the Pareto set of stable norm.

randomization, as it lies in the triangle (C, C_1, C_2). This construction gets us close to the desired norm s'.

However, one also needs to make sure that players are correctly incentivized to make a proposal, when they get an opportunity to do so, rather than to conceal this opportunity. This is the case if $\pi_1(C_1) \geq \pi_1(C)$ and $\pi_2(C_2) \geq \pi_2(C)$, i.e., if each player gets at least as high a payoff when he makes a proposal as when he remains silent. When one moves point C to C', these incentives might get violated, and the construction above must be adjusted as follows.

The new continuation payoff when no proposal is made, C', lies in between C and B. Suppose that it violates 1's incentives to make his prescribed proposal: $\pi_1(C') > \pi_1(C_1)$. Since, in the old norm, we had $\pi_1(C_1) \geq \pi_1(C)$, such a violation is possible only if $\pi_1(C_2) > \pi_1(C_1)$. In this case, we modify the prescribed proposal for player 1 by moving point C_1 towards C_2. As this happens, the value of $\pi_1(C_1)$ increases and the value $\pi_1(C')$ decreases (to keep B the weighted average). When these values become equal, the incentives for player 1 to make a proposal start holding again. With the new continuation payoff C_1' for player 1's proposal and renewed continuation payoff in case of no proposal C'', player 1 is incentivized to make the prescribed proposal. One then can check that both new points can be implemented: the payoff C_1' lies between C_1 and C_2 and therefore can be implemented using public randomization, while point C'' lies within the triangle (C, C_1, C_2) and can therefore also be implemented.

The same procedure is applied to player 2. The modified continuation payoffs can be implemented using public randomization device. The new norm s' therefore has the same set of payoffs as the old norm s at any stage, and it is self-sustaining at frequency q'.

I Non-vanishing probability of proposal

The sufficient conditions of Theorem 1 rely on the probability p of a player being able to make a proposal being proportional to ε: $p = q\varepsilon$. As explained at intuitively the beginning of Section 3, when p is independent of ε, one can no longer guarantee the existence of sustainable payoffs for all stage games. This section establishes the result formally: for $p \in (0, 1]$, there exists a stage game which has no sustainable payoffs when the discount factor δ is sufficiently close to 1.

An example of such a game is given by the matrix below, with some payoffs expressed in terms of a large constant M.

-2M,-2M	-2M,-2M	-2M,-2M	-1,7	-1,-2M	0,0
-2M,-2M	-2M,-2M	-2M,-2M	M,-2M	-2M,M	0,0
-2M,-2M	-2M,-2M	-2M,-2M	-2M,M	M,-2M	0,0
7,-1	-2M,M	M,-2M	-2M,-2M	-2M,-2M	0,0
-2M,-1	M,-2M	-2M,M	-2M,-2M	-2M,-2M	0,0
0,0	0,0	0,0	0,0	0,0	0,0

The minmax values of players are $\underline{v}_1 = \underline{v}_2 = 0$, as seen from the last row and the last column. The set of Pareto efficient payoffs is a part of a line which goes through points $(-1, 7)$ and $(7, -1)$, and is a segment between $(0, 6)$ and $(6, 0)$: any other stage game payoff gives strictly less total payoff of the players, $\pi_1 + \pi_2$.

LEMMA 2 *For any $p > 0$, there exists M such that for all ε small enough, no self-sustaining norm exists.*

Proof. The proof proceeds by contradiction. Suppose there exists a self-sustaining norm and let $\mathcal{V}$ denote the set of payoff vectors implemented by all continuations of the norm that follow histories at which no off-path proposal has been accepted. Consider the payoff vector $Q_2 \in \mathcal{V}$ that gives the highest payoff to player 2 in $\mathcal{V}$, and suppose that $\pi_2(Q_2) > 0$. If player 1 gets a chance to propose, 1 gets at least $\pi_1(Q_2)$. Therefore the infimum A_1 of 1's payoff over all elements of $\mathcal{V}$ satifies

$$A_1 \geq \varepsilon * \underline{v}_1 + (1 - \varepsilon)(p * \pi_1(Q_2) + (1 - p) * A_1)$$

because 1 is gets at least minmax payoff $\underline{v}_1$ during the action stage, and can guarantee himself a payoff $\pi_1(Q_2)$ when he makes a proposal.

When implementing Q_2, the expected per-period payoff of player 2 must be strictly positive, since Q_2 gives the maximal payoff to 2 in $\mathcal{V}$. As was shown above, player 2 has to play as a pure strategy one of the columns from the second to the fifth; and respectively, player 1 is able to deviate and get a payoff of at least M. Thus, 1's continuation payoff has to satisfy:

$$\pi_1(Q_2) \geq \varepsilon * M + (1 - \varepsilon)(p * \pi_1(Q_2) + (1 - p) * A_1).$$

Indeed, if player 1 deviates and gets M in the current period, during the renegotiation stage player 1 can secure the payoff $\pi_1(Q_2)$ if given a chance to propose, and otherwise gets at least A_1.

Combining the above inequalities yields

$$p(M - \pi_1(Q_2)) \leq (1 - p)(A_1 - \underline{v}_1),$$

which is impossible for M large enough. Therefore, there is no self-sustaining norm for which 2 gets a strictly positive payoff. By symmetry, the same holds for player 1. The only possible self-sustaining norm that remains is for both players to always minmax each other and get zero in each period. However, this norm is not robust to a proposal to move to a babbling norm with payoff $(3,3)$, which is implementable by the Folk Theorem as long as ε is small enough. ■

Although, existence may be an issue when p fixed, this need not be the case. In particular, if the per-period utility of a deviator is sufficiently low during the punishment phase, this will be suffice to deter deviations, even if the probability of renegotiation is large. There are many stage games for which there exist sustainable payoffs at all frequencies of renegotiation. An example of such a game, in which Pareto inefficient punishment is needed to support on path behavior, is given by the matrix below:

-14,0	0,0	0,6
-14,0	0,0	5,0
-14,1	2,1	0,0

The Pareto frontier of this game consists of all the payoffs which lie between the points $(0,6)$ and $(5,0)$. All these payoffs are sustainable as δ becomes sufficiently large, even if the value of p remains fixed.

Proposition 5 *There exists $\varepsilon_0 \in (0,1)$ such that $\forall \varepsilon < \varepsilon_0$, any payoff on the Pareto frontier of the game is sustainable for any value of $p \in [0,1]$.*

Proof. The proof of this proposition follows the construction of Appendix A.1. In order to implement Pareto-efficient payoff $Q = (0,6)$, player 2 has to choose the third column, while player 1 has to choose the first row. However, 1 may be tempted to choose the second row and get a payoff of 5. In order to implement Q, therefore, there must be a Pareto-inefficient punishment for player 1, sufficiently harsh to deter 1 from making this deviation, despite being able to propose a Pareto improvement with positive probability in each period. This is achieved as follows: when implementing Q, if player 1 deviates and chooses the second row instead of the first one, he gets an immediate benefit of 5ε. Player 1 is punished by moving continuation payoff to point $B = (-5.5\varepsilon, 1)$. At this stage, if player 1 gets a chance to renegotiate, he proposes point $C = (0.5\varepsilon, 1)$, otherwise the play moves to point $A = ([-5.5 - 6\frac{p/2}{1-p/2}]\varepsilon, 1)$, which gives player 1 a continuation payoff of at least -11.5ε. Point A is implemented by choosing the third row and the first column with a payoff $(-14,1)$, and at the renegotiation stage player 1 would have a continuation payoff of at least ε. With this construction, point Q deters any off-path proposals by player 1. ■

J Multiplayer agreements: Voter-dependent continuations and sequential voting

J.1 Voter-dependent continuations

Suppose, first, that continuation payoffs can depend arbitrarily on the voting decision of each player—except if everyone agrees on a proposal, in which case stability dictates that the proposal is implemented. With this high degree of flexibility, norms may be constructed so that all negotiation proposals are dissuaded and the Folk Theorem obtains.

Theorem 6 *For any feasible payoff vector π with $\pi_i > \underline{v_i}$ for all i, π is sustainable.*

To understand this result, we recall that in the underlying repeated game without negotiation, any strictly individually-rational payoff vector can be implemented for ε small enough by minmaxing any player i who deviates in actions, and switch to minmaxing any player $j \neq i$ who deviates when minmaxing player i. The same idea can be applied when negotiation is feasible, by deterring it as follows: if a player, i, proposes another continuation, everyone else is prescribed to reject the proposal and to start minmaxing player i. If another player, j, deviates from the prescribed rejection by accepting i's proposal, and all other players reject it, then players are prescribed to minmax j instead of i. If two or more players accept i's proposal, it is implemented, which guarantees that the norm satisfies our stability refinement. This prescription guarantees that it is always suboptimal for a player to unilaterally accept a proposal and, consequently, that it is also suboptimal to make any proposal. Unless some additional restrictions are imposed on the continuation payoffs, allowing for the possibility of renegotiation with three or more players thus has no more predictive power on the set of equilibria and payoffs than the standard Folk Theorem.

To prove Theorem 6 formally, observe that since $v \in \mathcal{F}$, the standard Folk Theorem implies that for ε small enough v can be achieved by an equilibrium of the underlying repeated game. This equilibrium can be embedded into a norm of the repeated game with renegotiation. According to this norm, no proposal is ever prescribed at any stage of the game. If a player i ever makes a proposal, other players are prescribed to reject it and the continuation payoff corresponds to punishing player i, as if i had deviated in action in the underlying equilibrium. If only one player $j \neq i$ accepts i's proposal, the continuation corresponds instead to the punishment equilibrium for j. If at least two players accept the proposal, it is implemented. These prescriptions guarantee that any unilateral deviation in action, proposal, or acceptance decision is suboptimal.

J.2 Sequential voting in case of no restrictions

Sequential voting permits more than two continuation payoffs, depending on the sequence of acceptance decisions of the players. The set of sustainable payoffs is qualitatively similar to the earlier analysis with simple norms but more permissive.

PROPOSITION 6 *Suppose that each proposal is decided by sequential voting. Then, sufficient and necessary conditions analogous to those of Theorem 4 obtain, which are characterized by upper orthants. Moreover, the set characterized by each of these two conditions is larger than the corresponding set obtained with simple norms and simultaneous voting.*

Sequential voting with many continuations thus provides more predictive power than simultaneous voting with voter-dependent continuations, but less predictive power than the simultaneous-voting specification with only two continuations.

Proof. [Sketch] Consider for simplicity the case of three players: player 1 makes a proposal and player 2 responds first, followed by player 3. Depending on responding players' votes, there are four possible continuations, one of which is equal to 1's proposal and arises when 2 and 3 accept the proposal.

The ability to punish 2 for accepting player 1's proposal is constrained by the following issue: if 2 accepts the proposal, 3 will reject it only if the punishment for player 2 gives him at least the same payoff as 1's proposal, which will be implemented if he accepts it. This puts a lower bound on 2's punishment payoff, which is higher than the minmax $\underline{v_2}$.

As a result, 1's punishment for making an off-path proposal is also limited. Since fewer punishments are available, fewer norms are self-sustaining: sequential voting has more predictive power than simultaneous voting. By nature of the arguments used to derive necessary and sufficient conditions, these conditions are characterized by upper orthants, even if players randomize their acceptance decision. Since allowing only two continuations—as simple norms do with simultaneous voting—is a special case of the more numerous continuations allowed by sequential voting, it follows that simple stable norms have more predictive power than the stable norms obtained with sequential voting. ■

K Observable mixed strategies

We have assumed throughout the paper that when a player randomizes across several actions or proposals, only the outcome of this randomization is observed by the other player. In particular, players' continuation values cannot directly depend on their choice of mixed strategy. Our results do not change if instead we assume that mixed strategies are observable. For sufficient conditions, this fact is straightforward because our construction is clearly compatible with players observing more information. For necessary conditions, payoff lower bounds were computed using only that any player can guarantee himself at least his minmax payoff during the action stage and at least some particular payoff during the proposal stage which satisfies the responder. These lower bounds do not change when mixing is observable.

The observability of mixed strategies does affect, however, the set of weakly renegotiation-proof (WRP) equilibria defined by Farrell and Maskin (1989), as follows. An SPE σ is *weakly-renegotiation proof* if there do not exist continuation equilibria σ^1, σ^2 of σ such that σ^1 strictly Pareto dominates σ^2. If a payoff vector arises as players' continuation payoff following some history of a WRP equilibrium, we will also say that these payoffs are WRP.

Assuming that mixing probabilities are observable, Farrell and Maskin obtained a sufficient condition for any feasible payoff to be WRP in the context of two-player repeated games. To formulate this condition, they define $c_i(\alpha) = max_{a'_i} \pi_i(a'_i, \alpha_{-i})$ as the *cheating* payoff of player i when he chooses a best response to the (mixed) action α_{-i}, and establish the following result.

PROPOSITION 7 *Let $\pi = (\pi_1, \pi_2)$ denote a feasible payoff. If there exist (mixed) action pairs $\alpha^i = (\alpha^i_1, \alpha^i_2)$ (for $i = 1, 2$) such that $c_i(\alpha^i) < \pi_i$, and $\pi_{-i}(\alpha^i) \geq \pi_{-i}$, then the payoff π is WRP if δ is sufficiently close to one.*

Moreover, with observable mixed strategies the set of WRP payoffs generically contains Pareto-efficient payoffs, as shown in Evans and Maskin (1989).

THEOREM 7 *Given the players' action spaces A_1 and A_2, for a generic choice of payoff functions, if players are sufficiently patient, then there exists a WRP equilibrium that is Pareto-efficient.*

We now prove the existence of a symmetric stage game in which all Pareto-efficient payoffs above the minmax satisfy the requirement of the above proposition, but cannot be WRP if mixing probabilities are unobserved, even if the stage game payoffs are slightly perturbed. The definition of WRP is the same as before, except that equilibrium strategies now depend only on the history of realized actions rather than on the history that included mixed strategies. The stage game is identical to the one described in Appendix I with $M = 100$, and is reproduced here for convenience:

-2M,-2M	-2M,-2M	-2M,-2M	-1,7	-1,-2M	0,0
-2M,-2M	-2M,-2M	-2M,-2M	M,-2M	-2M,M	0,0
-2M,-2M	-2M,-2M	-2M,-2M	-2M,M	M,-2M	0,0
7,-1	-2M,M	M,-2M	-2M,-2M	-2M,-2M	0,0
-2M,-1	M,-2M	-2M,M	-2M,-2M	-2M,-2M	0,0
0,0	0,0	0,0	0,0	0,0	0,0

As noted in Appendix I, the minmax values of players are $\underline{v}_1 = \underline{v}_2 = 0$, as seen from the last row and the last column. The set of Pareto efficient payoffs is a part of a line which goes through points $(-1, 7)$ and $(7, -1)$, and is a segment between $(0, 6)$ and $(6, 0)$: any other stage game payoff gives strictly less total payoff of the players, $\pi_1 + \pi_2$. Let's show that none of those Pareto efficient payoffs can be a part of WRP given low enough ε (arbitrarily patient players), even if players have access to a public randomization device. In fact, only the minmax payoff of $(0, 0)$ is WRP:

Counter-Example 1 *With unobservable mixed strategies,* $(0, 0)$ *is the unique WRP payoff.*

Suppose, by way of contradiction, that there is a point A that is the continuation payoff of some WRP equilibrium σ and such that $\pi_2(A) > 0$ (the case of $\pi_1(A) > 0$ is similar). Consider the payoff vector A' corresponding to player 1's lowest payoff and, hence, player 2's highest payoff among all continuation payoffs of σ before public randomization.[4] When implementing A', player 2 cannot choose either the first or the last column, since there column give him at most zero, regardless of 1's strategy; this would imply that 2's continuation payoff in the next period satisfies $\pi_2 > \pi_2(A')$, contradicting our choice of A'. Therefore, 2 chooses among columns located between the second and fifth.

Since A' gives 1 his lowest possible payoff, when implementing A' player 1 cannot get a period payoff higher than $\pi_1(A')$, even if he always plays a stage-game best response. Otherwise, the promise-keeping constraint would have to prescribe a continuation giving 1 a payoff lower than $\pi_1(A')$. If 2 chooses a pure strategy (among the columns from the second to the fifth), player 1 can guarantee himself a payoff of $M = 100$, which is greater than what 1 gets from any individually-rational payoff. Thus, when implementing A', player 2 should play a mixed strategy. Due to promise-keeping constraints, each pure action that 2 chooses in equilibrium should give 2 a payoff

[4] Since σ is WRP, 1's lowest continuation payoff is achieved for 2's highest continuation payoff. The proof can be easily adjusted if σ's payoff extrema are not achieved.

of at least $\pi_2(A') > 0$.[5] Such a strategy is impossible, however, as shown below.

When implementing A', let player 1 choose the rows from the first to the third with a total probability of β, and choose the fourth and the fifth rows with the total probability of γ. If player 2 chooses either the second or the third column, 2's expected payoff is at most $X_1 = -2M\beta + M\gamma$, while if 2 chooses either the fourth or the fifth column, 2's expected payoff is at most $X_2 = M\beta - 2M\gamma$. Since both payoffs X_1, X_2 cannot both be strictly positive, 2 has to mix either between the second and the third columns, or between the fourth and the fifth columns.

Let player 1 choose the fourth row with probability β', and the fifth row with probability γ'. If player 2 chooses the second column, 2's expected payoff is at most $X_1' = M\beta' - 2M\gamma'$, while if 2 chooses the third column, 2's expected payoff is at most $X_2' = -2M\beta' + M\gamma'$. Since both payoffs X_1', X_2' cannot be strictly positive at the same time, 2 cannot mix between the second and the third columns.

Let player 1 choose the first row with probability $\hat{\alpha}$, the second row with probability $\hat{\beta}$, and the third row with probability $\hat{\gamma}$. If 2 chooses the fourth column, his expected payoff is at most $\hat{X}_1 = 7\hat{\alpha} - 2M\hat{\beta} + M\hat{\gamma}$, while if he chooses the fifth column, his expected payoff is at most $\hat{X}_2 = -2M\hat{\alpha} + M\hat{\beta} - 2M\hat{\gamma}$. The sum of the two payoffs $\hat{X}_1 + \hat{X}_2 = (7 - 2M)\hat{\alpha} - M\hat{\beta} - M\hat{\gamma} \leq 0$. In particular, the payoffs $\hat{X}_1, \hat{X}_2$ cannot both be strictly positive and, hence, player 2 cannot be mixing between the fourth and the fifth columns.

Since there is no mixed strategy which gives player 2 a strictly positive payoff for each of 2's actions, point A is not WRP. The only WRP payoff is the minmax payoff $(0, 0)$. Moreover, even if one slightly perturbs the stage game payoffs, there will be no WRP equilibrium which provides the players with payoffs that are significantly higher than the minmax. By contrast, with observable mixed strategies and sufficiently patient players, one could implement any Pareto efficient payoff strictly above the minmax. Any Pareto efficient payoff satisfies Proposition 7 where α^1 corresponds to 1 choosing the fourth row and 2 mixing between the first and second columns with equal probability, and α^2 corresponds to 1 mixing between the first and second rows with equal probability and 2 choosing the fourth column.

[5] Indeed, 2's continuation payoffs are all below $\pi_2(A')$, by definition of A', while the pure actions in the support of his optimal mixed strategy give him an expected payoff of $\pi_2(A')$.

L Self-sustainable norms in the absence of proposer-specific punishments

One virtue of explicitly modeling the renegotiation process is to incorporate the logic of modern repeated games analysis into renegotiation: just as arbitrary continuation equilibria may follow from observed actions in a repeated game, here arbitrary continuations may follow rejected proposals. The paper has explored one consequence of this generality, which is that even good proposals may be deterred, and Pareto dominated equilibria be sustained as a result.

While in the applications discussed earlier this flexibility seemed reasonable or even desirable, in other environments it is natural to ask what equilibria may be sustained when proposers cannot be punished. Indeed, such a restriction is imposed in a number of models of explicit negotiation[6] and sometimes formalized as a "No-Fault Disagreement" (NFD) axiom. The axiom requires the continuation equilibrium following a rejected proposal to coincide with the default continuation in case no proposal was made. This appendix shows how our results are modified when this refinement is added.

In order to keep the language of the analysis as close as possible to the existing literature, this section adopts the "stable convention" terminology of Definitions 2 and 3 instead of the self-sustainability refinement.[7]

Definition 10 *A stable convention* $\mathcal{C}$ *is* forgiving *if for any equilibrium* s *in* $\mathcal{C}$*, for any* i *and* m_i*,* $s_{+1}[s](z, \boldsymbol{a}, i, m_i, 0){=}s_{+1}[s](z, \boldsymbol{a}, i, \emptyset, 0)$.

Our concepts are modified as follows. A payoff vector A is said to be *forgivingly q-sustainable* if for all ε small enough, there is a forgiving stable convention containing an equilibrium which expected payoff is equal to A. A is forgivingly sustainable if it is forgivingly q-sustainable for all q's large enough.

The main result in this case is given by the novel necessary conditions, which are much more restrictive those of Theorem 1: the continuation payoffs must lie within a distance $O(\frac{1}{q})$ of the convex hull of the (individually-rational) Pareto frontier. More precisely, for each feasible payoff vector A, let $\rho(A)$ denote the signed distance from the line (P_1P_2), counted positively if A lies below (P_1P_2), and negatively otherwise, as indicated by Figure 5.

[6] See Santos (2000) and Miller-Watson (2013). A similar idea appears in Farrell (1987), Rabin (1994), and Arvan, Cabral, Santos (1999) for the case of simultaneous announcements.

[7] The concepts used in this appendix can be readily re-expressed in terms of self-sustainability.

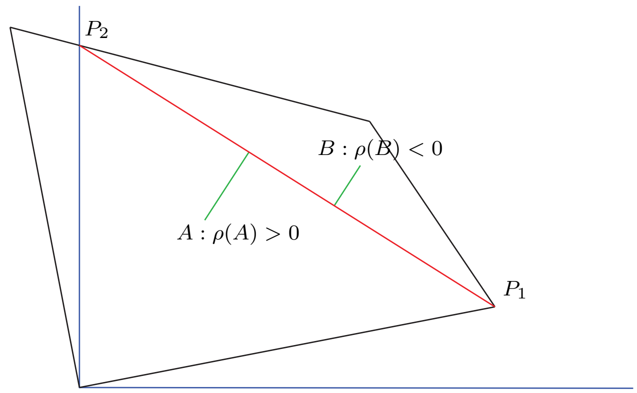

Figure 5: Signed distance from (P_1P_2)

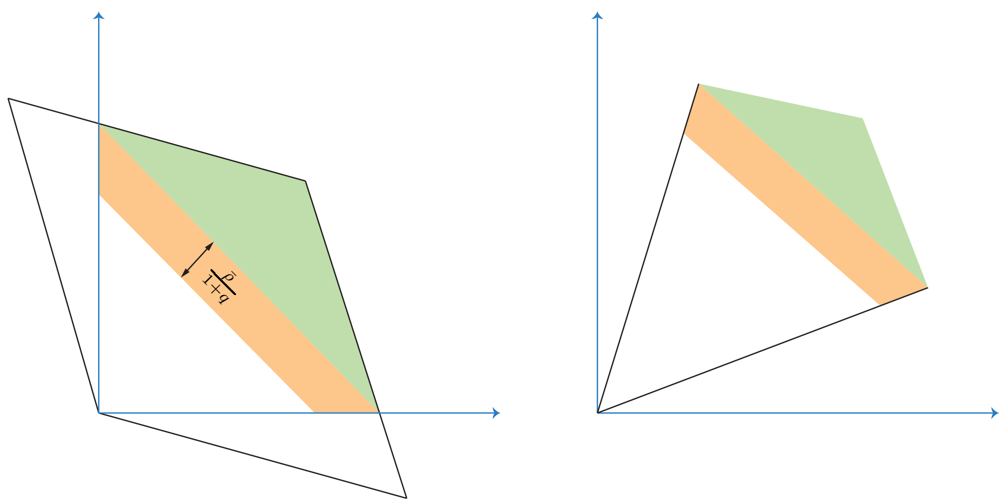

Figure 6: Sustainable payoffs without proposer-specific punishments

Let $\bar{\rho}$ denote the maximum value of ρ among all feasible payoff vectors.

Theorem 8 *If A is forgivingly q-sustainable, then $\rho(A) \leq \frac{\bar{\rho}}{1+q}$.*

One may also wonder whether all the feasible payoffs lying above the line (P_1P_2) can be achieved in this case. The next result provides a positive answer which is independent of negotiation frictions. To establish this result, we slightly modify the definition of stability, as follows: deviating proposal which is accepted needs to be implemented only if it improves the proposer's payoff by more than a constant $\eta > 0$, arbitrarily small but fixed, over his equilibrium payoff without the deviation.[8]

Definition 11 *A convention $\mathcal{C}$ is η-**stable** if a) at any element $s \in \mathcal{N}$ the players engage in rich renegotiation, and b) the following holds: consider any equilibrium of $\mathcal{C}$ and history at which i gets*

[8]Using the refinement in Theorem 8 affects the corresponding bound by a factor η.

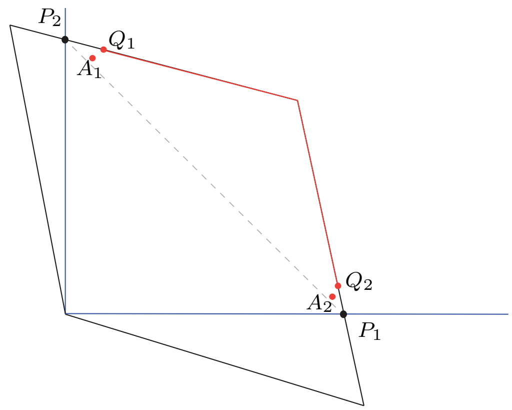

Figure 7: Construction of a convention with NFD (payoffs)

a chance to make a proposal and let $\hat{U}_i$ denote i's continuation payoff. Then, whenever i proposes a plan $s \in \mathcal{S}$ giving him at least $\hat{U}_i + \eta$, and $-i$ accepts it, μ is implemented.

Theorem 9 *Assuming η-stability, any payoff vector A strictly above the segment (P_1P_2) is forgivingly sustainable.*

The role of η is to prevent off-path proposals whose payoffs lie near the boundary of the convention's payoff set, as detailed in the proof of the theorem.

L.1 Proof of Theorem 9 (Sufficient Conditions)

Notation: throughout the analysis, for any payoff vector X achieved by some norm of $\mathcal{C}$, we will denote by $X^{\mathcal{C}}$ the corresponding equilibrium.

Consider two feasible Pareto points, Q_1 and Q_2, lying at an arbitrarily small but strictly positive distance from P_2 and P_1, respectively, and illustrated by Figure 7. It suffices enough to show that for any ε small enough, there exists a forgiving stable convention $\mathcal{C}$ which includes Q_1 and Q_1 as equilibrium payoffs, that is, convention has elements $Q_1^{\mathcal{C}}$, $Q_2^{\mathcal{C}}$. By public randomization, this will imply that this convention can also be made to contain all payoffs above the segment $[Q_1, Q_2]$. The argument below focuses on the case in which P_2 and P_1 are determined by the minmax payoffs, which is the harder one.[9]

We construct a convention which continuation payoffs just after the public randomization stage (before the action stage) consist of the Pareto frontier contained between Q_1 and Q_2 and of two

[9]If, say, $\pi_1(P_2) > \underline{v_1}$, it suffices to set $Q_1 = P_2$ in our construction and use it as as the best proposal for player 2.

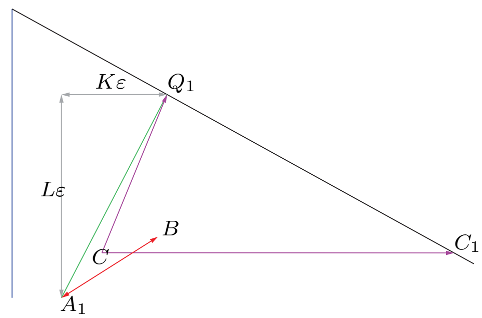

Figure 8: Construction of a convention with NFD (implementation)

additional points, A_1 and A_2, respectively lying within ε-proportional distance from Q_1 and Q_2, as indicated on Figure 7. We describe the implementation of $A_1^{\mathcal{C}}$ and $Q_1^{\mathcal{C}}$; $A_2^{\mathcal{C}}$ and $Q_2^{\mathcal{C}}$ have a symmetric implementation.

While Q_1 is taken as given, the location of A_1 depends on ε, and is determined by the following conditions

$$\pi_1(A_1) = \pi_1(Q_1) - K\varepsilon$$

$$\pi_2(A_1) = \pi_2(Q_1) - L\varepsilon, \tag{14}$$

for constants K and L which will be determined ulteriorly.

To implement $A_1^{\mathcal{C}}$, players are prescribed to minmax each other. The continuation payoff B after the action stage is a function of the players' realized actions, a_1 and a_2: $B = B(a_1, a_2)$. The implementation is illustrated by Figure 8. For any action a_i of player i the continuation payoff $\pi_i(B(a_i, a_j))$ does not depend on a_j.

Given that player 2 has minmaxed player 1, let $Eu_1(a_1)$ denote 1's expected payoff for the period, as a function of his chosen action, a_1. 1's continuation payoff, $\pi_1(B(a_1, a_2))$, satisfies the promise-keeping condition

$$\pi_1(A_1) = \varepsilon Eu_1(a_1) + (1 - \varepsilon)\pi_1(B(a_1, a_2)).$$

A similar relation holds for 2's continuation payoff. By appropriately choosing players' continuation payoffs $B(a_1, a_2)_{(a_1,a_2)\in\mathcal{A}}$, the construction can make players indifferent between taking *any* action in the game.

Moreover, if the constant K appearing in (14) is large enough, then for any action profile (a_1, a_2), one necessarily has $\pi_1(B(a_1, a_2)) < \pi_1(Q_1))$.[10]

[10] Indeed, the distance between A and $B(a_1, a_2)$ is proportional to ε, with a coefficient bounded above by the highest

Consider any of the continuation payoffs $B(a_1, a_2)_{(a_1,a_2)\in\mathcal{A}}$ after the action stage—henceforth referred to as 'B' for simplicity. B is a weighted average of three continuation payoffs corresponding to the following events: player 1 makes a proposal, player 2 makes a proposal, no one makes a proposal. Let C denote the continuation payoff in case no one makes a proposal (this payoff is computed before the public randomization taking place in the following period).

For the convention to be forgiving, any rejected proposal results in payoff C. This implies that if player 1 gets to make a proposal, in equilibrium he proposes the element with a Pareto-efficient payoff C_1 which gives 2 her default value $\pi_2(C)$, making player 2 to accept the proposal in equilibrium.

The situation is different if player 2 gets to make a proposal. $B^{\mathcal{C}}$ gives player 1 a lower payoff than $Q_1^{\mathcal{C}}$, and player 2 is prescribed to propose an element $Q_1^{\mathcal{C}}$, which achieves her highest payoff in the convention and also gives player 1 a higher payoff than $C^{\mathcal{C}}$ does.

As shown on Figure 8, at element $B^{\mathcal{C}}$ if player 1 gets a chance to make a proposal, he proposes $C_1^{\mathcal{C}}$, if 2 gets a chance to make a proposal, she proposes $Q_1^{\mathcal{C}}$. B is thus a weighted average of C, C_1 and Q_1. Given any point B, one can find a default option C such that B is indeed the right weighted average, given the probabilities of proposal for each player.

We will verify at the end of this proof that the constants K and L from (14) may be chosen so that C lies to the right of the line (A_1, Q_1). If this is true, $C^{\mathcal{C}}$ may be implemented, before public randomization, as a weighted average of $A_1^{\mathcal{C}}$, $Q_1^{\mathcal{C}}$, and $Q_2^{\mathcal{C}}$.

The remaining element of interest, $Q_1^{\mathcal{C}}$, is implemented as follows: players are prescribed to choose the pure-strategy Pareto-efficient payoff northwest of Q_1. If 1 deviates in action, the continuation payoff jumps to B; if 2 deviates, it jumps to the analog of B near Q_2. Players are incentivized to play as prescribed as long as $\frac{\pi_1(Q_1)-\pi_1(B)}{\varepsilon}$ is large enough. This is achieved by judiciously choosing the constants K and L arising in (14), as explained next.

Determination of the constants K and L

First, we observe that for K large enough, the threat of jumping to continuation $B^{\mathcal{C}}$ is enough to incentivize player 1 to play as prescribed in the implementation of $Q_1^{\mathcal{C}}$. We fix such a K—this choice is independent of ε. We now show that for L big enough, for any realization of B (which depends on which actions players choose while implementing $A^{\mathcal{C}}$), the point C will lie to the right of line A_1Q_1, as mentioned earlier.

absolute value of the payoff of the stage game.

Since a player's probability of proposal and the distance from B to the Pareto line are both proportional ε, the distance between B and C must be proportional to ε^2. Therefore, if we can show that each continuation point $B(a_1, a_2)$ lies to the right of the line A_1Q_1, at a strictly positive ε-proportional distance, so does the point C, for sufficiently small ε.

The points $B(a_1, a_2)$ are constructed by promise-keeping conditions. Let B^* denote the continuation payoff, out of all continuations $B(a_1, a_2)$, which gives the lowest payoff to player 1 and the highest payoff to player 2. B^* corresponds to the highest value $Eu_1(a_1)$ out of all actions a_1 and to the lowest value $Eu_1(a_2)$ out of all actions a_2. It suffices to show that B^* lies to the right of A_1Q_1. We recall the promise-keeping conditions

$$\pi_1(A_1) = \varepsilon Eu_1(a_1) + (1-\varepsilon)\pi_1(B^*)$$

$$\pi_2(A_1) = \varepsilon Eu_2(a_2) + (1-\varepsilon)\pi_2(B^*)$$

or, equivalently,

$$[\pi_1(A_1) - \pi_1(B^*)] = \varepsilon[Eu_1(a_1) - \pi_1(B^*)]$$

$$[\pi_2(A_1) - \pi_2(B^*)] = \varepsilon[Eu_2(a_2) - \pi_2(B^*)].$$

The ratio of the absolute values of the right-hand sides in the two equations above, $|\frac{Eu_2(a_2)-\pi_2(B^*)}{Eu_1(a_1)-\pi_1(B^*)}|$, determines the tangent of the angle of the vector A_1B^* above the horizontal. Since B^* is at an ε-distance from Q_1, this ratio simplifies to $|\frac{Eu_2(a_2)-\pi_2(Q_1)}{Eu_1(a_1)-\pi_1(Q_2)}|$, plus ε-terms which can be ignored.

Player 1 cannot obtain a higher payoff than his minmax $\underline{v}_1$ (as player 2 is minmaxing him), and player 2 cannot obtain a lower payoff than her lowest possible payoff in the game, which we denote as $\underline{v}$. Therefore, the angle of the vector A_1B^* above the horizontal is no higher than $|\frac{\underline{v}-\pi_2(Q_1)}{\underline{v}_1-\pi_1(Q_2)}|$, a finite value independent of L and ε.

The tangent of the angle of the line (A_1Q_1) above the horizonal is equal to $\frac{L}{K}$. By choosing L high enough, this ratio exceeds twice the ratio $|\frac{\underline{v}-\pi_2(Q_1)}{\underline{v}_1-\pi_1(Q_2)}|$. This guarantees that the vector A_1B^* lies strictly to the right of the line (A_1Q_1), as desired.

There remains to check that the convention satisfies all the conditions of Theorem 9. First, both players are incentivized to propose as prescribed: player 1 proposes the best available option for him, given the default option C. If player 2 wants to improve upon Q_1, she has to propose a continuation which gives her at least η more than her on-path continuation payoff. For ε small enough, however, the only proposals that would achieve this would have to give player 1 less than $\pi_1(C)$, and would therefore be rejected. Second, the continuation payoff, C, is the same when a proposal is rejected, regardless of the identity of the proposer and the nature of the proposal. The convention is thus forgiving. Finally, the point Q_1 is a continuation of the convention both after and before the public randomization, as desired.

L.2 Proof of Theorem 8 (Necessary Conditions)

Consider a forgiving stable convention $\mathcal{C}$. For simplicity, we assume that at each stage of the game—before the action stage, before the proposal stage, and before the public randomization stage—there exist equilibria in the convention with respective payoff vectors A, B, and C, that yield the maximal value of ρ at the corresponding stage.[11] Let α denote the (possibly mixed) action profile corresponding to the first-period play implementing element $A^{\mathcal{C}}$—the continuation before the action stage, and let $v(\alpha)$ denote the expected current payoff resulting from α. Since $\rho(v(\alpha)) \leq \bar{\rho}$, we necessarily have

$$\rho(A) \leq \varepsilon\bar{\rho} + (1-\varepsilon)\rho(B)$$

Point B, which is a continuation payoff before the proposal stage, is the weighted average of the continuation payoffs following accepted proposals, and of the default option. When a player—player 1, say—gets a chance to make a proposal, the expected continuation payoff must lie within at most an $\sqrt{\varepsilon}$-distance from the Pareto line. Otherwise, player 1 could propose a Pareto point which increases both players' payoffs by a value proportional to $\sqrt{\varepsilon}$, and is an equilibrium lying above the minmax.[12] This proposal would then be accepted by player 2 and would be a profitable deviation for player 1. Therefore, if a player gets a chance to make a proposal, which happens with probability $q\varepsilon$, the resulting continuation cannot have a positive value of ρ that exceeds $\sqrt{\varepsilon}$. When no one makes a proposal, the continuation payoff is dictated by the default continuation, whose value of ρ is at most $\rho(C)$. This implies that

$$\rho(B) \leq q\varepsilon \times \sqrt{\varepsilon} + (1-q\varepsilon)\rho(C).$$

Finally, since C is a convex combination of payoffs, obtained by public randomization, of equilibrium payoffs before the action stage whose maximal ρ-value is achieved by A,

$$\rho(C) \leq \rho(A).$$

Combining the above inequalities and getting rid of second-order ε-terms shows Theorem 8.

[11] If the supremum values are not achieved, the proof can be easily adjusted by taking appropriate limits.

[12] With the more permissive concept of an η-stable convention, the continuation payoff has to lie within a distance of $\sqrt{\varepsilon} + \eta$ from the Pareto line. Otherwise player 1 could make a proposal which gives him η more, and gives player 2 $\sqrt{\varepsilon}$ more than the continuation payoff.

L.3 Arbitrary number of players without proposer-specific punishment

Finally, consider the most restrictive case of a simple convention that is also forgiving, as defined in the two-player case.

Definition 12 *A simple convention $\mathcal{C}$ is* forgiving *if, for each period t and history $h \in \mathcal{H}^+$ ending in period t, the continuation play in case of a proposal not accepted by the supermajority is the same as if no proposal was made.*

The definitions of (forgivingly) sustainable payoffs at frequency q and (forgivingly) sustainable payoffs are identical to those of the two-player case.

The necessary conditions resemble the two-player case. Let $\mathcal{P}'$ denote the set of individually-rational Pareto-efficient payoffs and $Co(\mathcal{P}')$ denote the convex hull $\mathcal{P}'$.

Proposition 8 *If A is forgivingly q-sustainable, the distance from A to $Co(\mathcal{P}')$ is bounded above by a decreasing function of q, which converges to 0 as q becomes arbitrarily large.*

The proof closely mirrors the argument used for the two-player case and is only sketched here. Suppose that A is the point of the convention which has the largest distance from $Co(\mathcal{P}')$ and that A lies "too far" down away from $Co(\mathcal{P}')$. Whenever a player gets to make a proposal—which happens with probability proportional to q—he proposes a Pareto point (or close to it). Moreover, the continuation payoff A' which follows if the proposal is rejected cannot lie farther away from $Co(\mathcal{P}')$ than A does. Combining this puts a bound on A's distance to $Co(\mathcal{P}')$, which vanishes as q gets large.

We conclude this section with sufficient conditions.

Theorem 10 *Assuming η-stability, any point A in the set $Co(\mathcal{P}')$ lying strictly above the minmax is forgivingly sustainable.*

Proof. [Sketch] We construct a forgiving η-stable convention $\mathcal{C}$ as follows. The convention $\mathcal{C}$ includes all Pareto-efficient payoffs which lie at some arbitrary small, but ε-independent distance from the minmax values. The convention $\mathcal{C}$ also includes, for each player, a set of Pareto-inefficient elements used to build a punishment equilibrium for that player, all elements in each set lie within a distance of order ε from the Pareto-efficient elements of the convention. For each player i, there is a Pareto-inefficient payoff vector A_i which gives i his worst payoff in $\mathcal{C}$. The equilibrium $A_i^{\mathcal{C}}$ which achieves payoff A_i, together with its continuations, form the punishment set for player i, as described below.

If players were unable to make any proposal, one could implement payoff A_i as follows. Player i is being minmaxed, which may require other players to use mixed strategies. As described in earlier proofs, this results in a set $\mathcal{B}^1$ of continuation payoffs, (potentially) one for each observed action profile (these various continuations are needed to incentivize the minmaxing strategy). Each continuation payoff $B^1 \in \mathcal{B}^1$ is implemented by minmaxing player i, which again generates several continuation payoffs in the next period, with generic element denoted as B^2. Player i is minmaxed in this way for several periods. In each period i's continuation payoff, π_i, increases by an amount of order ε. One can compute the number T of periods needed to minmax player i, so that π_i exceeds $\pi_i(A_i)$ by a sufficiently high amount that i can be incentivized to play any action by the threat of returning to A_i. The value of T is independent of ε. After these T periods, each continuation payoff B^T can be implemented by playing a deterministic sequence of actions so that the continuation payoff always lies within some ε-proportional distance from B^T. This implementation is an equilibrium, since the payoff A_i prevents any deviation from player i, and any deviation by another player leads to an even larger drop in the continuation payoff of the deviator.

When proposals are re-introduced in the game, there will be changes in the implementation of $A_i^{\mathcal{C}}$, but these changes will be insignificant. After the first round of minmaxing player i, the resulting continuation payoff B^1 is calculated taken into account the possibility of proposals. That is, $B^{1\mathcal{C}}$ is the convex combination of some default option, $C^{1\mathcal{C}}$, if no one makes a proposal, and of proposals payoffs $C_i^{\mathcal{C}}$ for each player, which are chosen to be Pareto efficient elements of the convention $\mathcal{C}$. The distance between the payoffs B^1 and C^1 is of order ε^2—as explained the similar proofs seen earlier. In the next period, the continuation payoff before the actions will be C^1 (instead of B^1, in the previous paragraph). Therefore, if one repeats minmaxing player i for T periods, the resulting continuation payoff compared to the case with no proposals, will differ by an amount of order ε^2, which is negligible as ε becomes arbitrarily small. As the value of δ limits to one, the modified implementation of A_i, based on minmaxing player i for T periods and then choosing a deterministic sequence of actions, will thus be an equilibrium even with the possibility of proposals.

Finally, the payoff A_i (and, therefore, all the default continuation payoffs C's) can be chosen so as to lie within some distance $K\varepsilon$-distance from the Pareto line. With ε small enough, no player can make an off-equilibrium proposal that would give him a payoff of at least η more than the equilibrium proposal, while keeping all other players at least at well as off as with the default payoff C. Therefore, the constructed convention is η-stable. Using initial public randomization, one can then include in the convention any point in the convex hull $Co(\mathcal{P}')$, which concludes the proof. ∎

References

Arvan, L., Cabral, L., and V. Santos (1999) "Meaningful Cheap Talk Must Improve Equilibrium Payoffs," *Mathematical Social Sciences*, Vol. 37, pp. 97–106.

Evans, J., Maskin, E. (1989) "Efficient Renegotiation-proof Equilibria in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 361–369.

Farrell, J. (1987) "Cheap Talk, Communication, and Entry." *Rand Journal of Economics*, Vol. 18, pp. 34-39.

Farrell, J., Maskin, E. (1989) "Renegotiation in Repeated Games," *Games and Economic Behavior*, Vol. 1, pp. 327–360.

Mendelson, E (1997) *Introduction to Mathematical Logic*, 4th ed. Chapman & Hall, London.

Miller, D., Watson, J. (2013) "A Theory of Disagreement in Repeated Games with Bargaining," *Econometrica*, Vol. 81, pp. 2303–2350.

Rabin, M. (1994) "A Model of Pre-Game Communication," *Journal of Economic Theory*, Vol. 63, pp. 370-391.

Santos, V. (2000) "Alternating-Announcements Cheap Talk," *Journal of Economic Behavior & Organization*, Vol. 42, pp. 405–416.